Getting it Right
in an ON Demand World

(Course Code AM71)

Student Notebook

ERC 5.0

IBM Certified Course Material

Trademarks

IBM® is a registered trademark of International Business Machines Corporation.

The following are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

AIX
DS4000
eServer
HACMP
OpenPower
POWER5
POWER Hypervisor
RS/6000
System p5
Tivoli
xSeries
AIX 5L
DS6000
Enterprise Storage Server
IntelliStation
POWER
POWER5+
PowerPC
System i5
System Storage
TotalStorage
z9
BladeCenter
DS8000
Footprint
iSeries
POWER4
Power Architecture
pSeries
System p
System x
Virtualization Engine
zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX® is a registered trademark of The Open Group in the United States and other countries.

Linux® is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

December 2006 Edition

The information contained in this document has not been submitted to any formal IBM test and is distributed on an “as is” basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer’s ability to evaluate and integrate them into the customer’s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will result elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

© Copyright International Business Machines Corporation 2006. All rights reserved.
This document may not be reproduced in whole or in part without the prior written permission of IBM.
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Trademarks . . . . **xv**

Course Description . . . . **xvii**

Agenda . . . . **xix**

Unit 1. Introduction to the Course **1-1**
Welcome to the Course 1-2
Team IBM 1-3
Our Solutions Come in a Box 1-4
When we succeed… 1-5
When we fail… 1-6
Course Goals 1-7
Introductions 1-8
Certification Test 238 1-10
Course Agenda (1 of 2) 1-12
Course Agenda (2 of 2) 1-14
Any Questions So Far? 1-15

Unit 2. The Sales Cycle **2-1**
Unit Objectives 2-2
What is the Signature Selling Method (SSM)? 2-3
What is the Relevance of the SSM to this Course? 2-4
The Sales Cycle and the Buying Cycle 2-5
Stage 1 Noticing 2-6
Stage 2 Identifying 2-7
Stage 3 Validating 2-8
Stage 4 Qualifying 2-9
Stage 5 Conditionally Agreeing 2-10
Stage 6 Winning 2-11
Stage 7 Implementing 2-12
Checkpoint Questions (1 of 2) 2-13
Checkpoint Questions (2 of 2) 2-14
Unit Summary 2-15
References 2-16

Unit 3. Technology **3-1**
Unit Objectives 3-2
IBM and Technology 3-3
Technology and Innovation 3-4
When Technology Will Not ‘Fix' the Problem 3-5
Techno-babble, Buzzwords, and TLAs 3-6
Topic 1: System p Architecture 3-7
Topic Objectives 3-8

POWER Architecture 3-9
IBM's 2005 Patent Total: 13 Years of Leadership 3-10
IBM Power Architecture 3-11
IBM POWER Technology Roadmap for System p 3-12
POWER5+ Enhancements 3-13
POWER5+ Quad-Core Module Highlights 3-14
POWER5+ Packaging 3-15
POWER5 / 5+ Design 3-16
IBM System p5: Simultaneous Multi-Threading 3-17
POWER5 Performance 3-18
Checkpoint Questions (1 of 7) 3-19
Topic Summary 3-20
Topic 2: Reliability, Availability, and Serviceability (RAS) .3-21
Topic Objectives 3-22
Why is RAS Important to Our Clients? 3-23
What is the Autonomic Computing Initiative? 3-24
Enable High System Availability – Building on World-Class Hardware RAS 3-25
POWER5 RAS Additions 3-26
Checkpoint Questions (2 of 7) 3-27
Checkpoint Questions (3 of 7) 3-28
Topic Summary 3-29
Topic 3: Virtualization 3-30
Topic Objectives 3-31
Why IT Optimization and Virtualization? 3-32
What is System Utilization and Efficiency? 3-33
Virtualization 3-34
Value of Virtualization 3-35
IBM Virtualization Engine is the Technology and Services Framework to Help Clients Implement Virtualization 3-36
Advanced POWER Virtualization Option 3-37
LPAR and Dynamic LPAR 3-38
Shared Processor Partitions 3-39
Capped and Uncapped Partitions 3-40
Shared Processor LPAR - Responsive Partition Example 3-41
The Real Impact of Virtualization is Significant 3-42
Virtual I/O 3-43
Virtualization Simplifies Infrastructure 3-44
Virtual Networking 3-45
Virtual SCSI 3-46
Partition Load Manager 3-47
Integrated Virtualization Manager (IVM) 3-48
OS Support for Virtualization 3-49
The Value of Advanced POWER Virtualization 3-50
Scenarios – Where APV is Likely to Have Significant Impact 3-51
Take it to the Next Level… Ask Clients: 3-52
Quantifying TCO Benefits – EXAMPLE ONLY 3-53
Web Application Server Example 3-54
High Availability Example 3-55

Mixed Production and Test Environment Data Warehouse Customer Example . . 3-56
Mixed Workload - Server Consolidation Example 3-57
Resources 3-58
Checkpoint Questions (4 of 7) 3-59
Checkpoint Questions (5 of 7) 3-60
Topic Summary 3-61
Topic 4: Capacity on Demand 3-62
Topic Objectives 3-63
Why is Capacity on Demand Important to You? 3-64
CoD + Virtualization = System p Competitive Edge 3-65
What is Different in POWER5 Systems? 3-66
Permanent Capacity Upgrade on Demand (CUoD) 3-67
Reserve Capacity on Demand 3-68
On / Off Capacity on Demand 3-69
Trial Capacity on Demand 3-70
Sub-Capacity Licensing Guidelines* 3-71
Checkpoint Questions (6 of 7) 3-72
Topic Summary 3-73
Checkpoint Questions (7 of 7) 3-74
Unit Summary 3-75
Any Questions? 3-76

Unit 4. Servers 4-1
Unit Objectives 4-2
Introduction 4-3
Products and Technology 4-4
Selling the Technology 4-5
Buzzwords, Acronyms, and TLAs 4-6
Topic 1: IBM System p Servers 4-7
IBM Systems 4-8
Focus on IBM System p5 4-9
Why IBM System p5? 4-10
Evolution of the IBM System p5 Express Editions 4-11
IBM System p5 Express Editions 4-12
IBM System p5 Product Line 4-13
IBM System p5 185 Express (7037-A50) 4-14
IBM System p5 505 Express (9115-505) 4-15
IBM System p5 505Q Express (9115-505) 4-16
IBM System p5 510 Express (9110-51A) 4-17
IBM System p5 510Q Express (9110-51A) 4-18
IBM System p5 520 Express (9131-52A) 4-19
IBM System p5 520Q Express (9131-52A) 4-20
IBM System p5 550 Express (9133-55A) 4-21
IBM System p5 550Q Express (9133-55A) 4-22
IBM System p5 560Q Express (9116-561) 4-23
IBM System p5 570 (9117-570) 4-24
IBM System p5 575 (9118-575) 4-25
IBM System p5 590 (9119-590) 4-26

IBM System p5 595 (9119-595) 4-27
Topic 2: Hardware Management Consoles 4-28
Virtualization Management Options 4-29
Hardware Management Console (HMC) 4-30
Integrated Virtualization Manager (IVM) 4-31
IVM Limitations 4-32
Cross Platform Management with IBM Director 4-33
Topic 3: I/O Drawers and Racks 4-34
Expandable I/O Subsystems in pSeries 4-35
Remote I/O Architecture (RIO2) 4-36
I/O Drawer – Model 7311-D20 4-37
I/O Drawer – Model 7311-D11 4-38
24-inch I/O drawer (features 5791, 5794) 1 of 2 4-39
24-inch I/O drawer (features 5791, 5794) 2 of 2 4-40
IBM TotalStorage EXP24 (7031-D24, 7031-T24) 4-41
Server I/O Attachment IBM Sserver p5 Systems 4-42
Server I/O Attachment IBM System p5 4-43
IBM System p5 racks 4-44
Topic 4: Workstations 4-45
IBM IntelliStation POWER 185 4-46
IBM IntelliStation Power 285 4-47
Graphics Accelerators 4-48
Topic 5: IBM BladeCenter 4-49
Selling IBM BladeCenter 4-50
Optimized Workloads for BladeCenter 4-51
IBM BladeCenter 4-52
IBM BladeCenter JS20 4-53
IBM BladeCenter JS21 and JS21 Express 4-54
Checkpoint Questions (1 of 4) 4-55
Checkpoint Questions (2 of 4) 4-56
Checkpoint Questions (3 of 4) 4-57
Checkpoint Questions (4 of 4) 4-58
Unit Summary 4-59
Any Questions? 4-60

Unit 5. Infrastructure 5-1
Unit Objectives 5-2
Topic 1: Introduction 5-3
Introduction 5-4
What Do We Mean By Infrastructure? 5-5
The Value of Infrastructure 5-6
Topic 2: Networks 5-7
Topic Objectives 5-8
Networks Are Everywhere! 5-9
Advantages of System p Network Products 5-10
Buzzwords and Acronyms 5-11
Copper versus Fiber Optic 5-12
Network Adapters 5-13

Ethernet Adapter Choices 5-14
Integrated Ethernet Adapters 5-15
Single-port PCI-X Ethernet Adapters 5-16
Multiport PCI-X Ethernet Adapters 5-17
Asynchronous Transfer Mode (ATM) 5-18
Other Network Connections 5-19
Other Network Adapters 5-20
Other Network Adapters 5-21
Virtual Ethernet 5-22
Operating System Support 5-23
Checkpoint Questions (1 of 6) 5-24
Checkpoint Questions (2 of 6) 5-25
Topic Summary 5-26
Topic 3: Clusters 5-27
Topic Objectives 5-28
More Buzzwords! 5-29
Features of Cluster 1600 5-30
The Benefits of Clusters 5-31
Clusters Can Provide: 5-32
High Performance Computing (HPC) Clusters 5-33
Cluster for Simplified Management 5-34
Cluster 1600 Commercial Application Areas 5-35
Cluster 1600 Components 5-36
Cluster 1600 Nodes 5-37
Cluster 1600 Interconnects 5-38
IBM High Performance Switch (HPS) 5-39
Myrinet-2000 Switches 5-40
InfiniBand (IB) HCAs and Switches 5-41
Cluster 1600 Standard Interconnects 5-42
Cluster 1600 Software 5-43
Use Cluster System Management Software to: 5-44
Cluster Systems Management Choices 5-45
High Performance Software 5-46
GPFS – General Parallel File System 5-47
High Availability Clusters 5-48
HACMP 5.3 5-49
HACMP V5.4 5-50
HACMP for Linux V5.4 5-51
HACMP Extended Distance (XD) 5-52
Cluster 1600 Scalability Limits 5-53
Rules for Configuring Cluster 1600 5-54
Cluster 1600 Configuration Rules 5-55
Checkpoint Questions (3 of 6) 5-56
Checkpoint Questions (4 of 6) 5-57
Topic Summary 5-58
Topic 4: Grids 5-59
Topic Objectives 5-60
So, What is Grid Computing? 5-61

Grid Focus Areas and Industries 5-62
Why Should You Sell Grid Computing? 5-63
How Grid Drives Server Solutions 5-64
Where Can I Learn More About Grid? 5-65
Checkpoint Questions (5 of 6) 5-66
Checkpoint Questions (6 of 6) 5-67
Topic Summary 5-68
Unit Summary 5-69

Unit 6. Storage 6-1
Unit Objectives 6-2
The Storage Opportunity 6-3
Storage Challenge 6-4
Unique Synergy IBM System Storage and System p 6-5
The Benefits of this Synergy 6-6
Topic 1: IBM System Storage Disks 6-7
Benefits of IBM TotalStorage Disk Systems 6-8
Storage Buzzwords, Acronyms, and TLAs 6-9
IBM System Storage Disk Continuum 6-10
IBM System Storage EXP24 (7031-D24, 7031-T24) 6-11
PCI-X DDR Dual Channel Ultra 320 SCSI Adapters 6-12
Entry Level Storage Solutions 6-13
PCI-X 1Gbps iSCSI TOE Adapters 6-14
IBM System Storage N3000 – IP-Attached Storage 6-15
IBM System Storage N5000 – IP-Attached Storage 6-16
IBM System Storage DR550 and DR550 Express 6-17
IBM TotalStorage DS Family The industry's broadest range of disk storage systems 6-18
DS4000 Series 6-19
IBM System Storage External Disk Positioning 6-20
DS6000 Series 6-21
DS6000 Footprint Comparison 6-22
ESS 750/800 6-23
ESS EOM and New ESS Offering 6-24
DS8000 6-25
IBM System Storage External Disk Positioning 6-26
SAN Switches from SMB to Enterprise 6-27
SAN Volume Controller Storage Engine 6-28
Topic 2: Tape Storage 6-29
IBM TotalStorage LTO Tape Family 6-30
LTO Tape Drive Roadmap 6-31
IBM LTO Ultrium Offerings 6-32
TS3310 Tape Library 6-33
Statement of Direction 6-34
IBM TotalStorage Enterprise Tape Family 6-35
3584 Tape Library 6-36
3590 Tape Drive 6-37
3592 Tape Drive 6-38

IBM TotalStorage Tape Library/Server Offerings 6-39
Topic 3: Storage Software 6-40
IBM TotalStorage Open Software Storage Family 6-41
Disk Hardware Virtualization 6-42
SAN Volume Controller (SVC) V3.1 6-43
SAN Volume Controller (SVC) for Cisco MDS 9000 6-44
SAN File System (SFS) 6-45
IBM TotalStorage Productivity Center V3.1 6-46
Tivoli Storage Manager 6-47
Checkpoint Questions (1 of 2) 6-48
Checkpoint Questions (2 of 2) 6-49
Unit Summary 6-50
Any Questions? 6-51

Unit 7. Software 7-1
Unit Objectives 7-2
Topic 1: AIX5L 7-3
Topic Objectives 7-4
The Marketplace: Today and the Future 7-5
What Do Clients Want from an Operating System? 7-6
Why AIX5L? 7-7
Features of AIX5L V3 on Power5 7-8
AIX 5L™ V5.3 Technology Level 5 7-9
Checkpoint Questions (1 of 6) 7-10
Checkpoint Questions (2 of 6) 7-11
Topic Summary 7-12
Topic 2: Linux 7-13
Topic Objectives 7-14
The Operating Systems Battlefield 7-15
UNIX Revenue Share – Rolling 4 Quarters 7-16
What is Linux? 7-17
Why Linux on Power? 7-18
Linux and IBM System p Virtualization 7-19
IBM Linux on POWER Strategy 7-20
POWER Linux Discussion Points with Customers 7-21
Sources of Information 7-22
Checkpoint Questions (3 of 6) 7-23
Checkpoint Questions (4 of 6) 7-24
Topic Summary 7-25
Topic 3: IBM Software 7-26
Topic Objectives 7-27
What is Processor-Value Unit Licensing? 7-28
Per Processor Software Licensing 7-29
Fewer Processor Cores Are Needed Over Time 7-30
Software Vendors Differentiate by Processor Type 7-31
PVU Provides Flexibility through Granularity 7-32
Simple Migration of Existing Entitlements 7-33
New Processors Will Be Assigned Processor Value Units 7-34

Processor Value Unit Calculator 7-35
Topic Summary 7-36
Topic 4: Solutions 7-37
Topic Objectives 7-38
What Are Solutions? 7-39
Why Create Solution Editions and Solution Offerings? 7-40
Why Create Solution Editions and Solution Offerings? 7-41
IBM System p5 Solution Offerings 7-42
Introducing Solutions from System p 7-43
Why System p for Business Applications? 7-44
Why System p for Infrastructure / SOA? 7-45
Why System p for Databases? 7-46
Why System p for HPC? 7-47
IBM Server Consolidation Factory for System p 7-48
Opportunities for Server Consolidation 7-49
IBM Server Consolidation Factory for System p 7-50
Consulting and Deployment Services 7-51
IBM Tivoli Usage and Accounting Manager 7-52
Funding Solution Purchases 7-53
The Client Benefits of IT Financing 7-54
The Benefits of Choosing IBM Global Financing 7-55
Financing Solutions – ISV and IBM Seller Perspectives 7-56
Checkpoint Questions (5 of 6) 7-57
Checkpoint Questions (6 of 6) 7-58
Topic Summary 7-59
Unit Summary 7-60

Unit 8. Competition 8-1
Unit Objectives 8-2
Competitive Features of IBM System p 8-3
IBM Continues to Lead in UNIX Servers 8-4
Our Competitors Are Responding with New Offerings 8-5
Sun Opteron Product Roadmap for 2006 8-6
Sun X4600 Response – Attack 8-7
Sun Blade X8400 versus JS21 Blade with BladeCenter H 8-8
An Open Door 8-9
Continuing to Drive the IBM Systems Agenda 8-10
Two-pronged July Launch Strategy 8-11
System p5 Pricing Strategy – Attack and Defend 8-12
IBM BladeCenter JS2x Configuration 8-13
Migrating? Consider the Leader.. IBM 8-14
Performance and Benchmarks 8-15
Projected Benchmark Performance* 8-16
IBM Leads TOP500 with the World’s #1 Supercomputer BlueGene/L - US DOE/NNSA/LLNL – 280.6 TFLOPS 8-17
Notes on Benchmarks and Values 8-18
What is the Real Value of Performance to Clients? 8-19
POWER5 RAS Features 8-20

Virtualization in the Future 8-21
Virtualization in the Future 8-22
Virtualization in the Future 8-23
Available Partitioning Approaches 8-24
Partitioning Leadership 8-25
Business Value of Advanced POWER Virtualization 8-26
Positioning the System p Servers 8-27
Positioning the System p Servers 8-28
Positioning the System p Servers 8-29
p5-570 Positioning 8-30
p5-570 Express Price Competitiveness 8-31
System p5 575:8-core Excellent HPC Performance and Scalability 8-32
System p5 575:16-core Excellent HPC Performance and Scalability 8-33
IBM Data Warehousing Balanced Configuration Unit 8-34
p5-595: Consolidation on a Whole New Level* 8-35
Checkpoint Questions (1 of 2) 8-36
Checkpoint Questions (2 of 2) 8-37
Unit Summary 8-38

Unit 9. Sizing the Solution 9-1
Unit Objectives 9-2
Sizing and Solution Design 9-3
Sizing Concepts 9-5
The Sizing Model 9-6
Sizing, Resizing, and Capacity Planning 9-7
The Sizing Problem 9-8
Sizing Levels 9-9
Sizing Levels 9-10
Sizing Levels 9-11
Sizing is a Multi-dimensional Problem 9-12
Topic 1: Practical Sizing 9-13
A Practical Sizing Method 9-14
What is the Type of Application? 9-15
What is the User Doing (Workload Type)? 9-16
What is the Type of User Interface? 9-17
How Much Data is Being Processed? 9-18
What is the Number of Users? 9-19
What is the Work Rate and Transaction Size? 9-20
Performance Calculations 9-22
Rules of Thumb 9-24
Topic 2: Sizing and System Performance 9-25
Balanced System Performance 9-26
Utilization: The Good, the Bad, and the Ugly 9-27
Utilization (1 of 2) 9-28
Utilization (2 of 2) 9-29
RDBMS Rules of Thumb 9-30
RDBMS Data and File System Cache (1 of 2) 9-31
RDBMS Data and File System Cache (2 of 2) 9-32

RDBMS Raw Data to Disk Ratio . . . 9-33
Topic 3: The Sizing Process . . . 9-34
Sizing and Assumptions . . . 9-35
Sizing Tasks . . . 9-36
Sizing Inputs . . . 9-37
Sizing Outputs . . . 9-38
Who Performs Sizing? . . . 9-39
What Do the Sizing Experts Have in Common? . . . 9-41
Sizing and Resizing Process – Sizing Tasks . . . 9-42
Sizing and Resizing Process – Tuning Tasks . . . 9-43
Sizing and Resizing Process – Resizing Tasks . . . 9-44
Components Involved in Sizing . . . 9-45
Hardware Requirements . . . 9-46
Hardware Requirements - Example . . . 9-47
Building Block Choices . . . 9-48
Weighing Sizing Components . . . 9-49
Topic 4: Benchmarks . . . 9-51
Benchmarks . . . 9-52
Benchmarks and Solution Sizing . . . 9-53
OLTP Benchmarks . . . 9-55
Business Intelligence Benchmarks . . . 9-57
e-business Benchmarks . . . 9-59
SPEC JBB2000 Benchmark . . . 9-60
SPECweb99 Benchmark . . . 9-61
High Performance Computing Benchmarks . . . 9-63
LINPACK Benchmark . . . 9-65
ISV Benchmarks . . . 9-67
Oracle Applications Standard Benchmark . . . 9-68
Siebel Platform Sizing and Performance Program . . . 9-70
Relative Performance Benchmark . . . 9-71
Topic 5: The Balanced System Guideline . . . 9-73
The Balanced System Guideline . . . 9-74
Spreadsheet Pros and Cons . . . 9-76
General Sizing . . . 9-78
Correct Processor Configuration . . . 9-79
Balanced Systems . . . 9-80
CPU Magic Number Calculations . . . 9-81
CPU Magic Number Calculations . . . 9-82
Estimating CPU Power . . . 9-83
Estimating Memory Size . . . 9-84
Estimating Disk Size . . . 9-85
Balanced System Guideline - Introduction . . . 9-86
Balanced System Guideline - Balanced . . . 9-87
Balanced System Guideline - Balanced (1 of 2) . . . 9-88
Balanced System Guideline - Balanced (2 of 2) . . . 9-90
Balanced System Guideline - LPAR . . . 9-92
Balanced System Guideline – LPAR Planning . . . 9-93
Balanced System Guideline – Disk Planning . . . 9-94

Balanced System Guideline - Pricing 9-95
Balanced System Guideline - Price/Performance 9-96
Balanced System Guideline – Sizing CPU 9-97
Balanced System Guideline – Stats 9-98
Balanced System Guideline – Sizing Disks 9-99
Balanced System Guideline – Sizing Results 9-100
Balanced System Guideline – Calibration 9-101
Balanced System Guideline – Resize CPU 9-102
Balanced System Guideline – Resize RAM 9-103
Balanced System Guideline – Resize Disks 9-104
Balanced System Guideline – Resize Disk Use 9-105
Balanced System Guideline – Tx Modelling 9-106
Topic 6: Application Specific Sizing 9-107
Application-Specific Sizing 9-108
AIX Tools for Data Gathering 9-109
e-Config and Solution Design 9-110
System Planning Tool (SPT) 9-111
SPT Installation 9-112
Checkpoint Questions (1 of 2) 9-113
Checkpoint Questions (2 of 2) 9-114
Any Questions? 9-115
Unit Summary 9-116

Unit 10. Validating the Solution 10-1
Unit Objectives 10-2
What is a Solution Assurance Review (SAR)? 10-3
What Are the Benefits of a SAR? 10-5
When Should a SAR Be Done? 10-7
Pre-sale SAR 10-8
Pre-installation SAR 10-9
What Type of SAR Should I Perform? 10-10
When is an Expert SAR Mandatory? 10-11
How Do I Request a Solution Assurance Review? 10-12
The Client Environment and Expectations 10-14
Pre-Sales Checklists 10-15
Pre-Install Checklists 10-16
Product Specific Checklists 10-17
Checkpoint Questions (1 of 2) 10-18
Checkpoint Questions (2 of 2) 10-19
Any Questions? 10-20
Unit Summary 10-21
References 10-22
Let’s take a break 10-23

Appendix A. Checkpoint Solutions A-1

Trademarks

The reader should recognize that the following terms, which appear in the content of this training document, are official trademarks of IBM or other companies:

IBM® is a registered trademark of International Business Machines Corporation.

The following are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

AIX®
AIX 5L™
BladeCenter®
DS4000™
DS6000™
DS8000™
eServer™
Enterprise Storage Server®
Footprint®
HACMP™
IntelliStation®
iSeries™
OpenPower™
POWER™
POWER4™
POWER5™
POWER5+™
Power Architecture™
POWER Hypervisor™
PowerPC®
pSeries®
RS/6000®
System i5™
System p™
System p5™
System Storage™
System x™
Tivoli®
TotalStorage®
Virtualization Engine™
xSeries®
z9™
zSeries®

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX® is a registered trademark of The Open Group in the United States and other countries.

Linux® is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Course Description

Getting it Right in an ON Demand World

Duration: 3 days

Purpose

This course is designed to provide students with up to date information about IBM System p products and services that will enhance their ability to design and provide solutions for clients.

Audience

The intended audience for this course is primarily those people who need to consult, design, and implement technical solutions for clients which include IBM System p products.

Prerequisites

To get the most benefit from this course the student should have a minimum of six months experience with IBM System p products in a sales or technical pre-sales support role.

Objectives

After completing this course, you should be able to:

- Understand the overall methodology of designing a balanced solution
- Establish credibility with the customer through product knowledge, solutions knowledge, and use of support tools and resources
- Collect and analyze customer requirements
- Determine a suitable solution
- Understand capabilities and limitations of technology and their impact on solution design and be able to explain them to the customer
- Provide technical assistance to sales and account teams

Agenda

Day 1

Unit 1 - Introduction to the Course
Unit 2 - Sales Cycle
Unit 3 - Technology
Unit 4 - Servers

Day 2

Review Day 1
Unit 5 - Infrastructure
Unit 6 - Storage
Unit 7 - Software
Unit 8 - Competition

Day 3

Review Day 2
Unit 9 - Sizing the Solution
Unit 10 - Validating the Solution
Class Review

Unit 1. Introduction to the Course

What This Unit Is About

This unit introduces the students to the course and provides an opportunity for the students to introduce themselves to the instructor and the class. This unit covers course goals, agenda and schedule.

What You Should Be Able to Do

After completing this unit, you should be able to:

- List the goals of the course
- Understand the course agenda

How You Will Check Your Progress

Accountability:

- Ask questions if you have any.

Welcome to the Course

Before we start…

- Start/stop times
- Lunch
- Breaks
- Facilities
- Certification testing

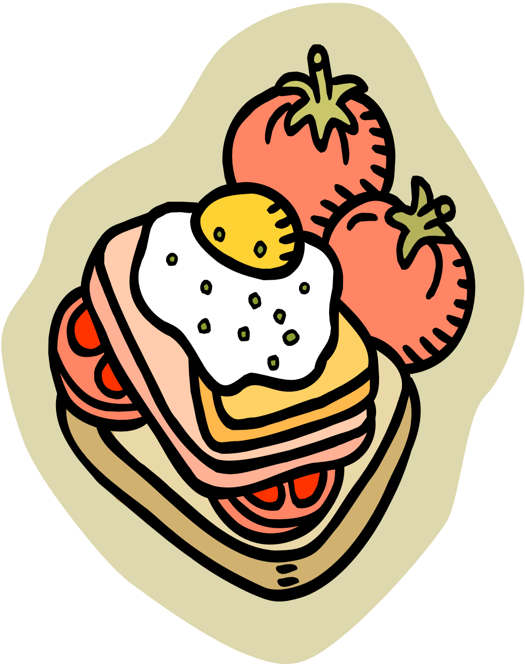

Figure 1-1. Welcome to the Course AM715.0

Notes:

Welcome to the Getting it Right in an On Demand World class. Before you start the instructor will tell you about the local arrangements for this class.

Team IBM

Sells solutions NOT boxes!

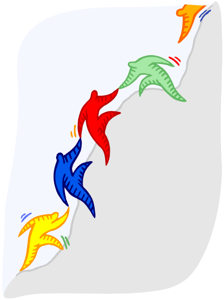

But…

Figure 1-2. Team IBM AM715.0

Notes:

We often hear how important it is that we sell ‘solutions' to our clients. The underlying idea being that customers are not all that interested in the technology as much as they are in the result, the business benefit. In my experience it depends on who in the customer organization you are talking about. I know lots of system administrators who are very much interested in the technology that will be delivered as part of the solution. But one fact is inescapable, when the solution arrives at the customer’s site it generally involves the delivery of lots of boxes!

Our Solutions Come in a Box

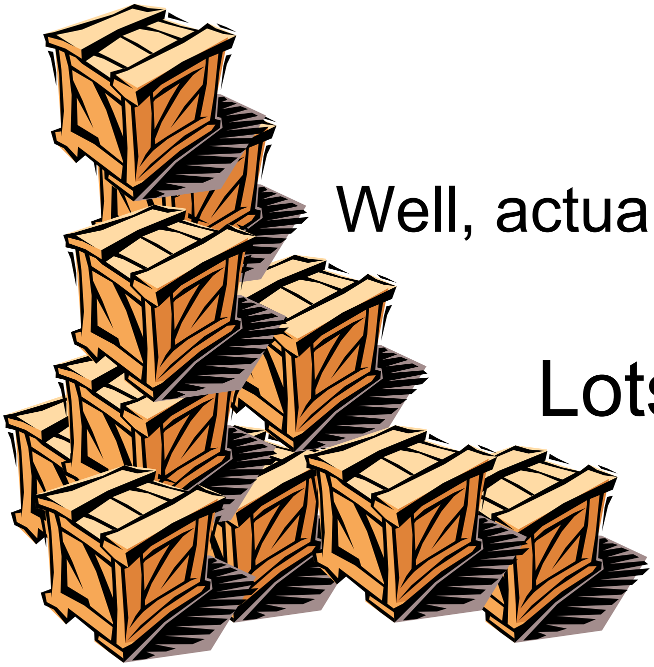

Well, actually…

Lots of boxes.

Figure 1-3. Our Solutions Come in a Box AM715.0

Notes:

The contents of the boxes and the skill which consultants like you bring to the customer should be able to bring about the ‘solution’ that the customer desires. Part of the work of this class is to provide information about the technologies and products that can help you craft more successful solutions for customers.

When we succeed…

then = Solution

and

Figure 1-4. When we succeed… AM715.0

Notes:

Simply put, if what's in the box yields what the client thinks is a solution then you have succeeded. If not….

When we fail…

Then = “critsits”

Figure 1-5. When we fail… AM715.0

Notes:

The result of missing expectations and failing to deliver the solution that the client had expected generates problems for you in the short term and could affect your chances of working with this client again. Critsits are an acknowledgement that sometimes we get it wrong. Reducing the number of critsits is a worthy goal, not just because it means the client is happier with the solution we provide but also because we have a greater chance of working with them in the future and developing a trusting consultative relationship that will benefit both the client and Team IBM.

Course Goals

- To provide IBM System p and Business Partner pre-sales technical support staff with information and tools that will help them identify, define and deliver the best solutions to our clients.
- To provide information for people planning to take the IBM certification test 238.

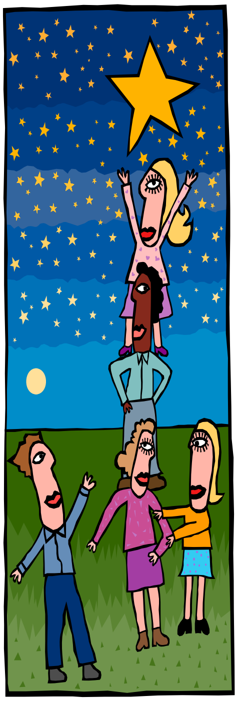

Figure 1-6. Course Goals AM715.0

Notes:

The principle goal of this course is to help you to deliver better solutions to your clients. In the next three days you will receive lots of valuable information that you can use to identify, define and create better solutions. The course also contains pointers to further information which you can explore to enhance your understanding. This course is a recognized learning resource for technical sales people who wish to achieve certification as *IBM eServer Certified Specialist - p5 and pSeries Technical Sales Support*. Preparation for the certification test is a great way to sharpen your skills and passing the test provides evidence of your achievement. This course, coupled with your experience and study should assist you in your preparation.

Introductions

- What is your current role?
- How long have you been doing that job?
- What do you want to learn from this class?

Figure 1-7. Introductions AM715.0

Notes:

Introductions

This is your opportunity to tell the rest of the class what kind of work you do and how long you have been doing the job. Tell the class and the instructor what you want to learn from the class; if you have specific objectives state them now.

Course Audience

Ask yourself, does this describe your role?

- Works with customers to gather customer requirements.
- Represents IBM eServer p5 and pSeries technologies to the customer's technical team.
- Takes a consultative approach to support sales.
- Trains customers and mentors peers in all aspects of the p5 and pSeries server line.
- Has detailed technical knowledge about the IBM eServer p5 and pSeries products, technologies, and solutions.

- Has skills to architect solutions to meet customer requirements.
- Performs competitive analyses.
- Translates technical requirements into customer solutions.
- Solves technical problems
- Describes p5 and pSeries strategy and solutions, and business trends and directions.

Course Prerequisites

Do you meet the course prerequisites?

- At least three years of experience in the Information Technology industry
- At least one year experience working with the IBM pSeries brand
- Knowledgeable in IBM practices and procedures, including tools, contacts, and resources

Are you in the right class?

If your role is very different from the one described above or if you do not feel that you meet the pre-requisites please discuss this with the instructor. It may well be the case that your personal objectives in taking the course can be met and that this course could be beneficial for someone in your role, but it is best to have the conversation now rather than feel that you wasted your time later.

Certification Test 238

- About the Test
  - 75 questions
  - 90 minutes duration
  - 61% passing score
- Preparation for the Test
  - Gain experience
  - Take this course
  - Take the assessment test
- Test Registration
  - Thomson Prometric
  - Pearson VUE

Figure 1-8. Certification Test 238 AM715.0

Notes:

About the Test

Certification test 238 consists of 70 multiple-choice questions which you must answer in a ninety minute session. To achieve certification you must score at least 61 points on the test.

Test Objectives:

Details of Test Objectives can be found at:
http://www-03.ibm.com/certify/tests/obj238.shtml for IBMers or at
http://www-03.ibm.com/certify/partner/mem/europe/en/ for Business Partners.

Prepare for the Certification Test

Test 238 uses an online assessment solution, AssessTrac which can be accessed at http://www.testrac.com/rs6000a/.

There is a charge ($10) for taking the assessment test online. The online assessment consists of seventy multiple-choice questions to be answered in ninety minutes and has a similar format to the certification test. The ability to go forward and backward through the test and to review and change answers is also provided. On completion of the assessment you will receive your results score by objective area as a percentage and with an indication of Pass or Fail for each section and for the whole test.

Register for the Certification Test

The test is taken at a Testing Center and requires that you register and reserve a time for the certification test. There are two companies which can provide this: Pearson VUE http://www.vue.com/ and Thomson Prometric http://securereg3.prometric.com/. Both companies require that you register with them before being able to schedule testing. If you already have taken another certification test with IBM you should use your Candidate ID to ensure that any certifications you take will be correctly added to the certification database.

Take the Certification Test

If this is your first certification test then I would recommend using AssessTrac first to get some experience not only with the type of questions that are asked as well as taking a test online. The results of the test will be calculated while you are at the site so you will know before you leave whether you have passed or not. You will be given an Examination Score Report which will serve as proof of your success and a formal certificate, well actually a ‘pdf' file will be sent to you some time later. Record you Candidate ID for future reference.

Course Agenda (1 of 2)

- Day 1
  - Unit 1 Introduction
  - Unit 2 Sales Cycle
  - Unit 3 Technology
- Lunch
  - Unit 3 Technology
  - Unit 4 Servers

- Day 2
  - Review Day I
  - Unit 5 Infrastructure
  - Unit 6 Storage
- Lunch
  - Unit 7 Software
  - Unit 8 Competition

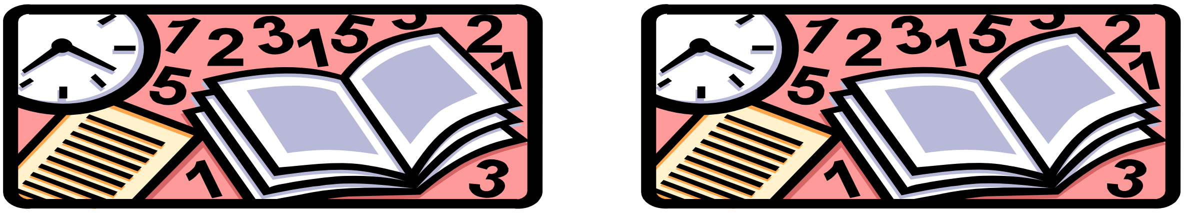

Figure 1-9. Course Agenda (1 of 2) AM715.0

Notes:

Course Agenda Day 1 Day 2

It is always a good idea to know what lies ahead. In the next three days this course will provide lots of information. Let's have a quick look at the order of presentation.

Day 1

Day 1 starts with an overview of the Sales process (The Sales Cycle) as described in the Signature Selling Method. The sales cycle is presented as an organizing principle for the course. This section is followed by detailed information about pSeries technologies that can provide you with the means to solve customer business problems (The Technology Parts 1 and 2). Understanding the way these technologies are packaged is the focus of the next section (The Products). The products section is divided into three parts: Servers, Storage, and Software. The Products Section starts on Day 1and concludes on Day 2.

Day 2

Its useful to review regularly so we start Day 2 with a quick review of Day 1 material and answer any questions before moving on to the final part of the products section. After lunch on the second day we will be looking at practical methods for sizing pSeries solutions. With your knowledge of the technologies and the products this should bring the process of solution design into focus. The remaining four teaching hours of Day 2 are devoted to the Sizing process.

Course Agenda (2 of 2)

- Day 3
  - Review Day 2
  - Unit 9 Sizing the Solution
  - Unit 10 Validating the Solution
- Lunch
  - Class Review
  - Questions, Certificates, Evaluations

Figure 1-10. Course Agenda (2 of 2) AM715.0

Notes:

Course Agenda Day 3

On Day 3 following the review of Day 2's material we will spend two hours on validating the solution. This section deals primarily with the Solution Assurance process. Next we turn to Delivering the Solution and will discuss the preparation for delivery, installation and commissioning of the solution. To conclude the course there will be a review test which will provide you with an opportunity to see how much of the information you have retained and areas to focus on in your preparation for certification test 238.

Certificates and Evaluations

At the conclusion of the course certificates of attendance will be provided and course evaluations will be completed by the students and instructor. The method for completing course evaluations may vary depending on the location of the class.

Any Questions So Far?

Figure 1-11. Any Questions So Far? AM715.0

Notes:

Questions

If you have questions now is the time to ask. Well, actually any time is a good time to ask. This concludes the introduction to the course.

Unit 2. The Sales Cycle

What This Unit Is About

This unit describes the Signature Selling Method sales cycle.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Explain the benefits of Signature Selling Method (SSM)
- List the seven stages of the selling cycle
- List the equivalent seven stages of the buying cycle
- Identify for each stage in the cycle:
  - Key activities
  - Measurable outcomes
  - Sales tools

How You Will Check Your Progress

Accountability:

- Checkpoint

References

http://w3-03.ibm.com/sales/compass/ssm_home/ssm_overview1.html

http://w3-03.ibm.com/sales/compass/ssm_home/ssm_train.html

Unit Objectives

On completion of this unit you should be able to:

- Explain the benefits of the Signature Selling Method (SSM).
- List the seven stages of the selling cycle.
- List the equivalent seven stages of the buying cycle.
- Identify for each stage in the cycle:
  - Key activities
  - Measurable outcomes
  - Sales tools

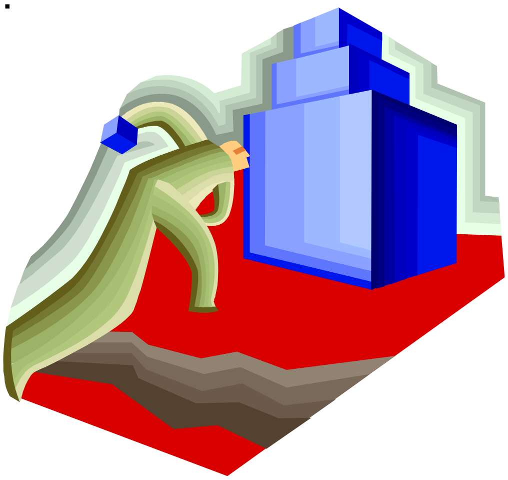

Figure 2-1. Unit Objectives

AM715.0

Notes:

What is the Signature Selling Method (SSM) ?

- SSM is the way Team IBM plans and makes customer sales calls to identify, develop, and close opportunities.
- It defines what Team IBM does with customers.
- Selling is more than what you say or do with a customer.
- Selling also includes planning and management.

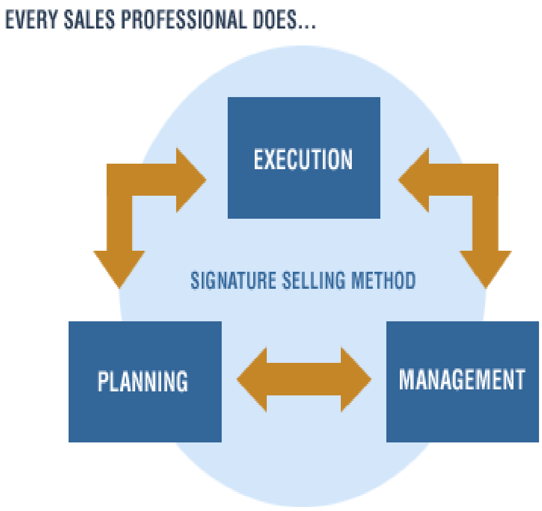

Figure 2-2. What is the Signature Selling Method (SSM)? AM715.0

Notes:

Planning

During the sales process, sellers ask these planning questions:

Where is the customer in the buying process?

Who is the right person to call on?

What do we need to do to assist the customer and advance the sale?

Execution

Execution is what you say and do with the customer to add value and to advance the sale.

Management

Successful sellers also self-manage their opportunities. They track the status of opportunities and periodically take snapshots of their opportunity pipeline to determine what corrective action, if any, is required.

What is the Relevance of the SSM to this Course?

- SSM provides:
  - Common language
  - Common outcomes
  - Common definition
- Which should:
  - Provide basis for shared model
  - Improve communications

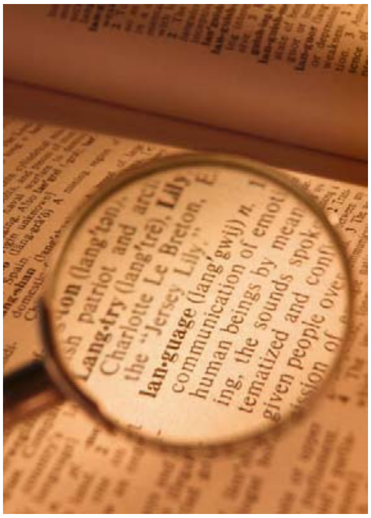

Figure 2-3. What is the Relevance of the SSM to this Course? AM715.0

Notes:

SSM has a common language, common outcomes, and common definitions of how to move through the selling cycle.

Team IBM can use the SSM to focus on the best solutions for customers, to speak the same sales language, to communicate better with other IBMers and Business Partners, and to help Team IBM act as one cohesive sales force.

The Sales Cycle and the Buying Cycle

Selling

1. Noticing
2. Identifying
3. Validating
4. Qualifying
5. Conditionally Agreeing
6. Winning
7. Implementing

Buying

1. Evaluating Strategies
2. Refining Strategies
3. Recognizing Needs
4. Evaluating Options
5. Selecting the Solution
6. Resolving and Deciding
7. Implementing the Solution and Evaluating Success

Figure 2-4. The Sales Cycle and the Buying Cycle AM715.0

Notes:

Because the SSM is customer-focused, each SSM stage corresponds to a customer buying stage. Even though the SSM has seven stages and begins with Stage 1, it is important to note that as sellers, IBMers and Business Partners start the SSM process at the point where the customer is in the buying process. Team IBM must be aligned with the customer.

Stage 1 Noticing

- Discuss business issues
- Build relationship
- Ask questions such as:
  - Where is the customer today?
  - Where would the customer prefer to be, in relation to strategic goals, competitive performance, financial performance, and other factors?

Figure 2-5. Stage 1 Noticing AM715.0

Notes:

At this stage, the customer and Team IBM are not focused on IT solutions; they are focused on business issues, core business processes (for example, supply chain), and strategies to make the processes run more efficiently and more effectively. Team IBM's experience with other customers can help create possibilities for improving business processes that the customer does not see. These possibilities are called thought leadership opportunities.

Stage 2 Identifying

- Aligning with potential decision makers
- Assessing the opportunity
- Determining why the customer needs the solution
- Determining if there are additional opportunities

Figure 2-6. Stage 2 Identifying AM715.0

Notes:

SSM Stage 2 begins when Team IBM selects and pursues one of the potential opportunities identified in Stage 1. The remainder of the SSM stages follows the life of this opportunity as Team IBM develops, closes, and implements it.

In Stage 2, Team IBM explores with the customer one of the issues or challenges that is creating a need for investment and action. As in Stage 1, Team IBM continues to focus on the customer's issues and needs because that is what the customer cares about and values as the basis for a relationship, whether long-term or project-specific. In Stage 2, Team IBM assesses whether the customer is serious about taking action on the issue immediately.

Stage 3 Validating

- Bridging from the customer's stated initiatives to envisioning the business capabilities that will solve the problem
- Anticipating how a Team IBM solution will meet the customer's need better than any other option
- Influencing the customer's buying options by suggesting capabilities that the customer might not have considered

Figure 2-7. Stage 3 Validating AM715.0

Notes:

In SSM Stage 3, the customer is focused on a specific business issue that is driving a need for action. The customer might think of this action as a project. For example, increasing inventory costs might prompt projects to optimize the supply chain, such as a project to improve demand forecasting or production planning. Team IBM wants to add value in this stage by helping the customer identify new or improved business capabilities wanted or needed by the customer in the enabling solution for the project, such as having real-time inventory status.

The customer is not focused on an information technology (IT) solution at this point. However, Team IBM needs to anticipate how a Team IBM solution would fit the required capabilities. During SSM Stage 3, the team completes the following activities:

Bridging from the customer's stated initiatives to envisioning the business capabilities that will solve the problem

Anticipating how a Team IBM solution will meet the customer's need better than any other option

Stage 4 Qualifying

- Review and influence the customer's buying and evaluation criteria.
- Develop and articulate Team IBM's initial solution.
- Evaluate the concerns of the key decision leader.
- Evaluate the risk to Team IBM and disengage if appropriate.
- Determine the level of presale Solution Assurance Review needed.

Figure 2-8. Stage 4 Qualifying AM715.0

Notes:

The objective of SSM Stage 4 is to gain the support of the key decision leader for the initial Team IBM solution. For the first time during the sale, the customer and Team IBM shift focus from the problem to the solution.

Stage 5 Conditionally Agreeing

- Moving to a final proposal
- Determining if the solution meets the customer's expectations
- Refining the value proposition
- Validating competitive strategy; Conducting the appropriate level of presale Solution Assurance Review

Figure 2-9. Stage 5 Conditionally Agreeing AM715.0

Notes:

During SSM Stage 5, Team IBM works with the customer to develop the proposed solution. The customer, who has been evaluating options, is ready to make a selection. Although the sale is not formally completed until Stage 6, Team IBM's goal in Stage 5 is to gain the key decision leader's conditional approval of the proposed solution.

Stage 6 Winning

- Signing a contract or statement of work
- Resolving any concerns
- Planning for change
- Making necessary refinements to the solution and negotiating final details

Figure 2-10. Stage 6 Winning AM715.0

Notes:

SSM Stage 6 ends with the customer signing the contract. Any issues and questions that might stall the sale should have been discovered and resolved prior to this stage of the sales process. There should be no surprises from unresolved problems at this stage if the SSM has been followed properly and the verifiable outcomes achieved for each previous stage.

Stage 7 Implementing

- Measure benefits
- Review progress
- Manage customer expectations
- Extend value
- Create opportunities
- Create reference story

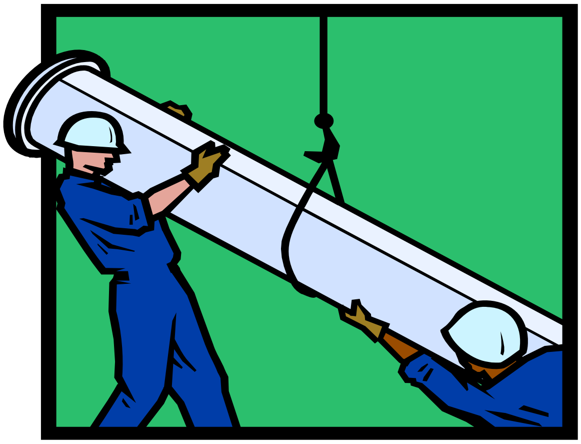

Figure 2-11. Stage 7 Implementing AM715.0

Notes:

During SSM Stage 6, the customer signed the contract, but this does not end Team IBM's sales efforts. During SSM Stage 7, Team IBM prepares to monitor the implementation and to ensure that the customer's expectations are met.

From the perspective of time, effort, and resources, developing opportunities within an established and successful account relationship is easier than selling to a new account. Much of the information collected in the previous SSM stages, such as customer business initiatives and relationship building, can be used to pursue additional opportunities.

The ability to leverage previous successes and to draw on established relationships can occur only if opportunities have been implemented successfully. Without a successful implementation, it can be difficult, if not impossible, to secure future business or sell-on opportunities, especially if Team IBM has to sell to the same contacts in the same buying decision unit. A successful implementation is also necessary if Team IBM is going to use the customer and solution to create reference stories.

In the final stage of SSM, Team IBM works to monitor the implementation and to ensure that the customer's expectations are met.

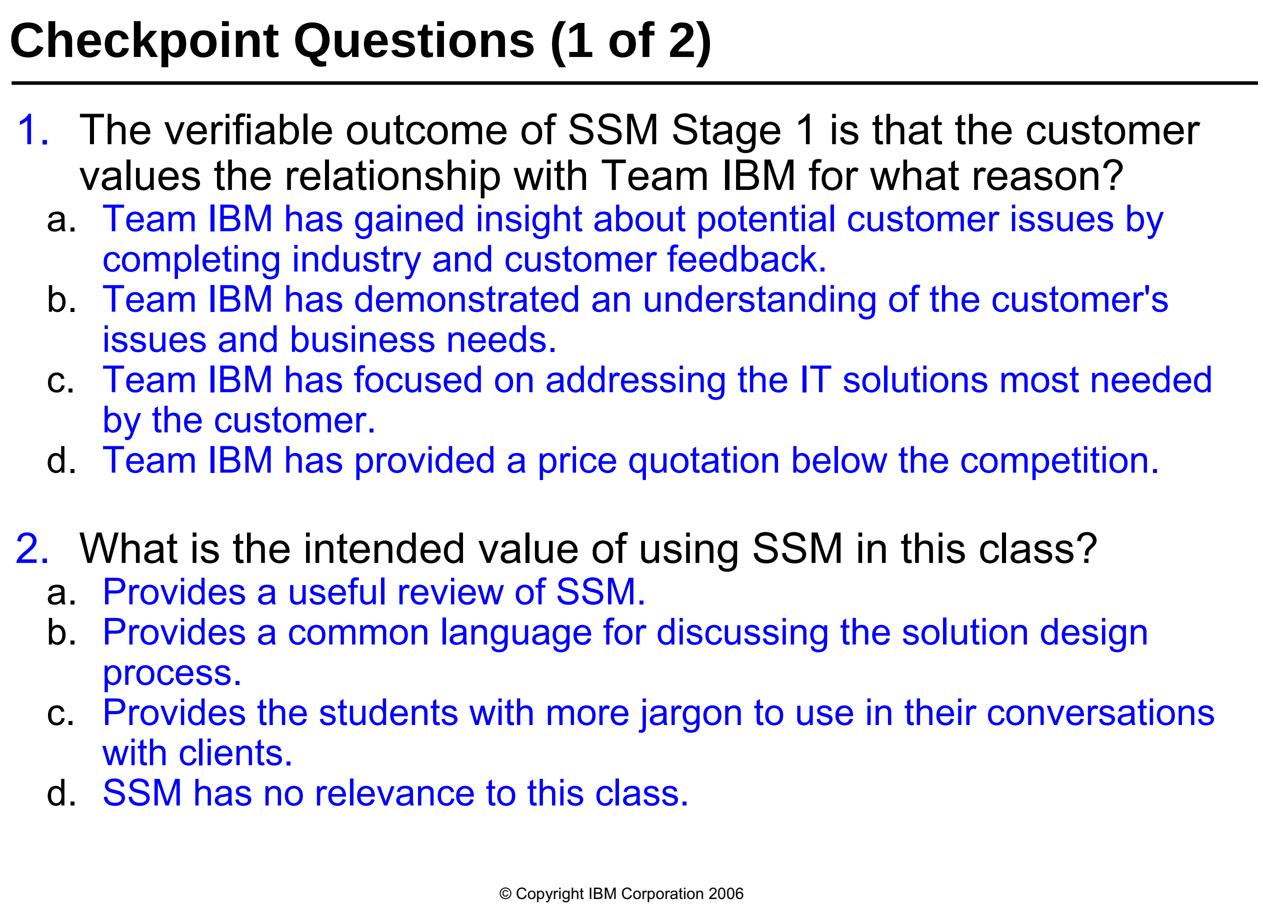

Checkpoint Questions (1 of 2)

1. The verifiable outcome of SSM Stage 1 is that the customer values the relationship with Team IBM for what reason?
   a. Team IBM has gained insight about potential customer issues by completing industry and customer feedback.
   b. Team IBM has demonstrated an understanding of the customer's issues and business needs.
   c. Team IBM has focused on addressing the IT solutions most needed by the customer.
   d. Team IBM has provided a price quotation below the competition.

2. What is the intended value of using SSM in this class?
   a. Provides a useful review of SSM.
   b. Provides a common language for discussing the solution design process.
   c. Provides the students with more jargon to use in their conversations with clients.
   d. SSM has no relevance to this class.

Figure 2-12. Checkpoint Questions (1 of 2) AM715.0

Notes:

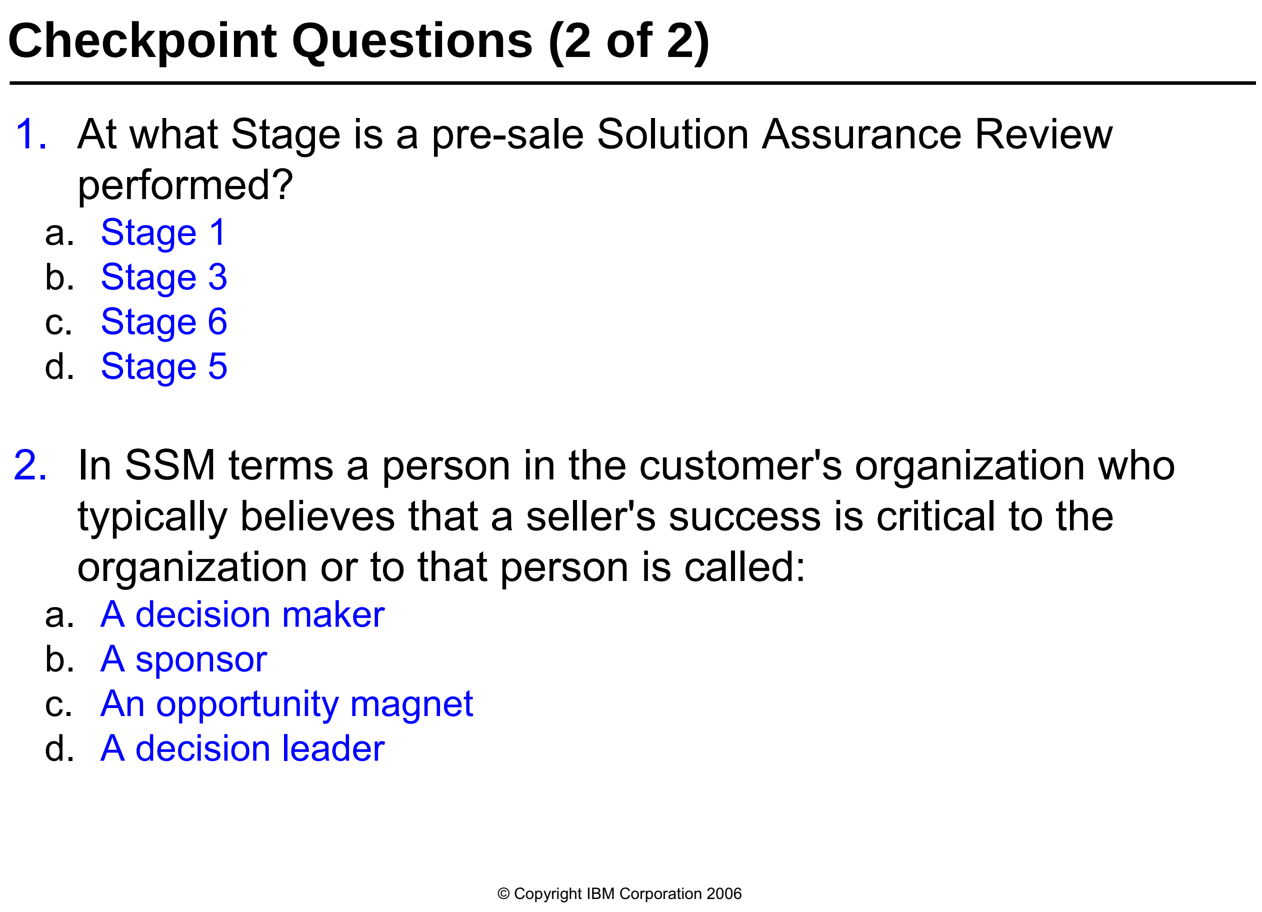

Checkpoint Questions (2 of 2)

1. At what Stage is a pre-sale Solution Assurance Review performed?
   a. Stage 1
   b. Stage 3
   c. Stage 6
   d. Stage 5

2. In SSM terms a person in the customer's organization who typically believes that a seller's success is critical to the organization or to that person is called:
   a. A decision maker
   b. A sponsor
   c. An opportunity magnet
   d. A decision leader

Figure 2-13. Checkpoint Questions (2 of 2) AM715.0

Notes:

Unit Summary

- You should now be able to:
  - Explain the benefits of the Signature Selling Method (SSM).
  - List the seven stages of the selling cycle.
  - List the equivalent seven stages of the buying cycle.
  - Identify for each stage in the cycle:
    - Key activities
    - Measurable outcomes
    - Sales tools

Figure 2-14. Unit Summary AM715.0

Notes:

References

- You may find more information here:

http://w3-03.ibm.com/sales/compass/ssm_home/ssm_overview1.html

http://w3-03.ibm.com/sales/compass/ssm_home/ssm_train.html

Figure 2-15. References AM715.0

Notes:

Unit 3. Technology

What This Unit Is About

This unit describes the technologies which are part of the IBM System p products.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Describe the relationship between technology and solutions
- List key IBM technologies that are part of the Power5 products
- Describe the functional benefits that these technologies provide
- Discuss the appropriate use of these technologies

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

After completing this unit you should be able to:

- Describe the relationship between technology and solutions.
- List key IBM technologies that are part of the Power5 products.
- Describe the functional benefits that these technologies provide.
- Discuss the appropriate use of these technologies.

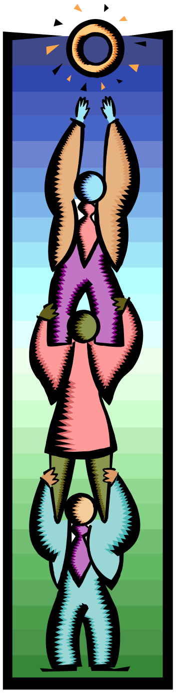

Figure 3-1. Unit Objectives AM711.0

Notes:

IBM and Technology

Solutions

Products

Technology

Science

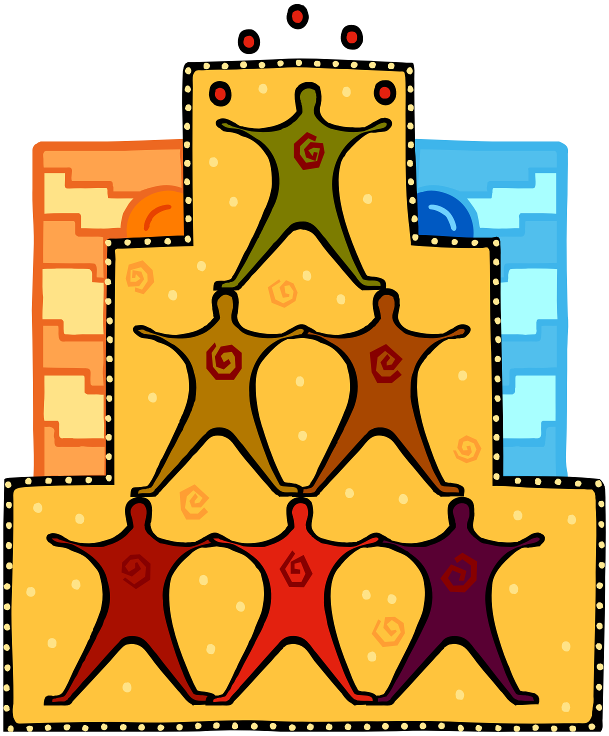

Figure 3-2. IBM and Technology AM711.0

Notes:

The pursuit of scientific discovery provides the basis for new technologies which can be incorporated into new and better products which can then enable clients to solve business problems.

This section is divided into two parts and will take approximately two hours to complete.

IBM's rich history of discovery and innovation has brought international recognition. In addition to five Nobel prizes, IBM researchers have been recognized with five U.S. National Medals of Technology, five National Medals of Science and 19 memberships in the National Academy of Sciences. IBM Research has more than 46 members of the National Academy of Engineering and well over 300 industry organization fellows.

Over the years, we have received international recognition for our discoveries and produced 22,357 patents - nearly 7,000 more than the nearest competitor. What is more important than the statistics is the effect these discoveries and patents are having in the marketplace -- and that's what really makes something innovative. Our continued innovation springs from our creative, dedicated people whose work continues to shape the future for our customers, the IT industry and the world.

Technology and Innovation

- Having technology available is a necessary first step
- Finding creative new ways to use the technology for the benefit of our clients is what innovation is about
- Solution design is an opportunity for innovative application of technology
- Innovation is the point at which invention meets insight.

Figure 3-3. Technology and Innovation AM711.0

Notes:

What do you think of when you think of innovation? Can you provide examples from your own experience? How would you define innovation?

When Technology Will Not 'Fix' the Problem

- When the technology is not related to the problem.
- When the client has unreasonable expectations.

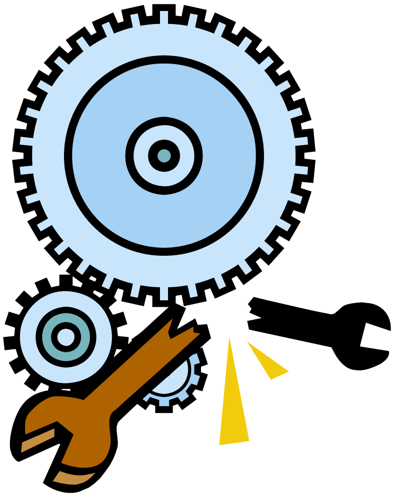

Figure 3-4. When Technology Will Not 'Fix' the Problem AM711.0

Notes:

That advances in technology can be applied to problems confronting our clients is not an issue, this is what we do! However, consider the case where the technology that we provide fails to solve the problem to the client's satisfaction. When this happens what might be one of the possible causes of dissatisfaction? It has been demonstrated that the degree of benefit you get from applying a particular technology is directly related to its appropriateness in the situation. Amdahl's Law shows this relationship.

Secondly the client may have unreasonable expectations. Setting expectations is certainly part of a successful solution design process but for the purpose of this section we will focus more on the technologies that are available and consider what problems might be solved by them. We will also look at some of the possible misapplications of the technology and their consequences.

Can you think of examples in your own experience where expectations were not met by the technology that was provided? Was the reason a failure of the technology, the expectations of the client or both?

Techno-babble, Buzzwords, and TLAs

- Talking about technology often leads to speaking in acronyms.
- In this unit you will be exposed to lots of them.
- Acronyms can provide us with useful short cuts in communication.
  - We say 'scuzzy' for SCSI for small computer system interface
- Techno-babble:
  - Speaking in acronyms without understanding the technology behind them.
  - Creates confusion and is to be avoided.

Figure 3-5. Techno-babble, Buzzwords, and TLAs AM711.0

Notes:

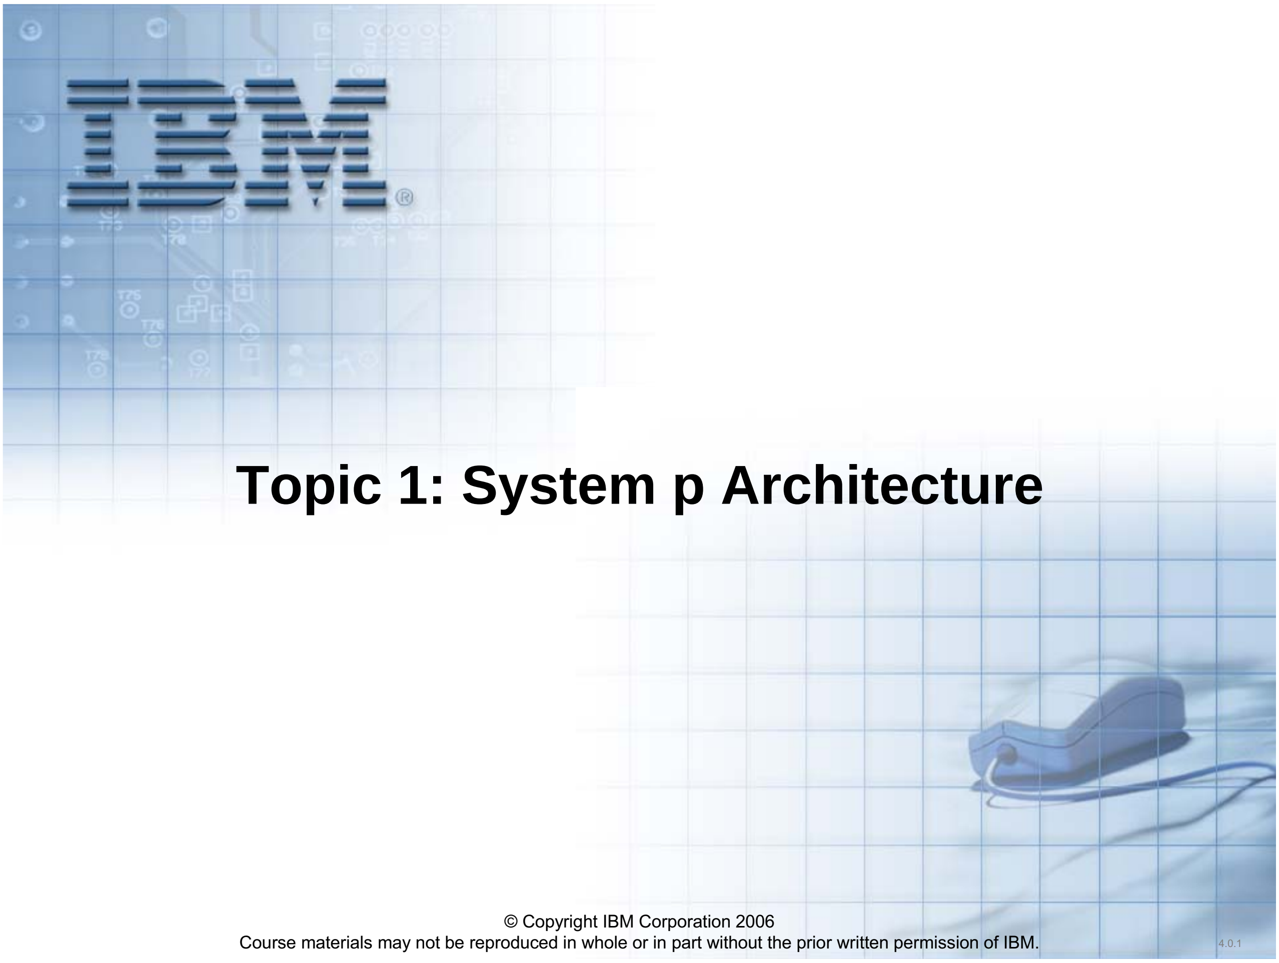

Topic 1: System p Architecture

4.0.1

Figure 3-6. Topic 1: System p Architecture AM711.0

Notes:

Topic Objectives

At the end of this topic you should be able to:

- Describe the advantages of RISC processor architecture.
- Explain the features of the POWER microprocessor family.
- Describe the processor roadmap for System p.
- Describe key enhancements in POWER5.

Figure 3-7. Topic Objectives AM711.0

Notes:

POWER Architecture

- 1980 Reduced Instruction Set Computer (RISC)
  - Superscalar
- 1990 POWER system
  - **P**erformance **O**ptimization **W**ith Enhanced **R**ISC
- Currently in fifth generation POWER5

John Cocke 1925 -2002

Figure 3-8. POWER Architecture AM711.0

Notes:

John Cocke (May 30, 1925 - July 16, 2002) was an American computer scientist recognized for his large contribution to computer architecture and optimizing compiler design. He is considered by many to be “the father of RISC architecture.”

1980:

IBM builds first prototype computer employing RISC (Reduced Instruction Set Computer) architecture. Based on an invention by IBM scientist John Cocke in the early 1970's, the RISC concept simplified the instructions given to run computers, making them faster and more powerful. Today, RISC architecture is the basis of most workstations and UNIX servers and is widely viewed as the dominant computing architecture of the future.

1990:

IBM announces its new RISC System-based computer line, the RS/6000 (now called IBM eServer pSeries), running AIX V3. Architecture of the system is given the name POWER (POWER1), standing for Performance Optimization With Enhanced RISC.

IBM’s 2005 Patent Total: 13 Years of Leadership

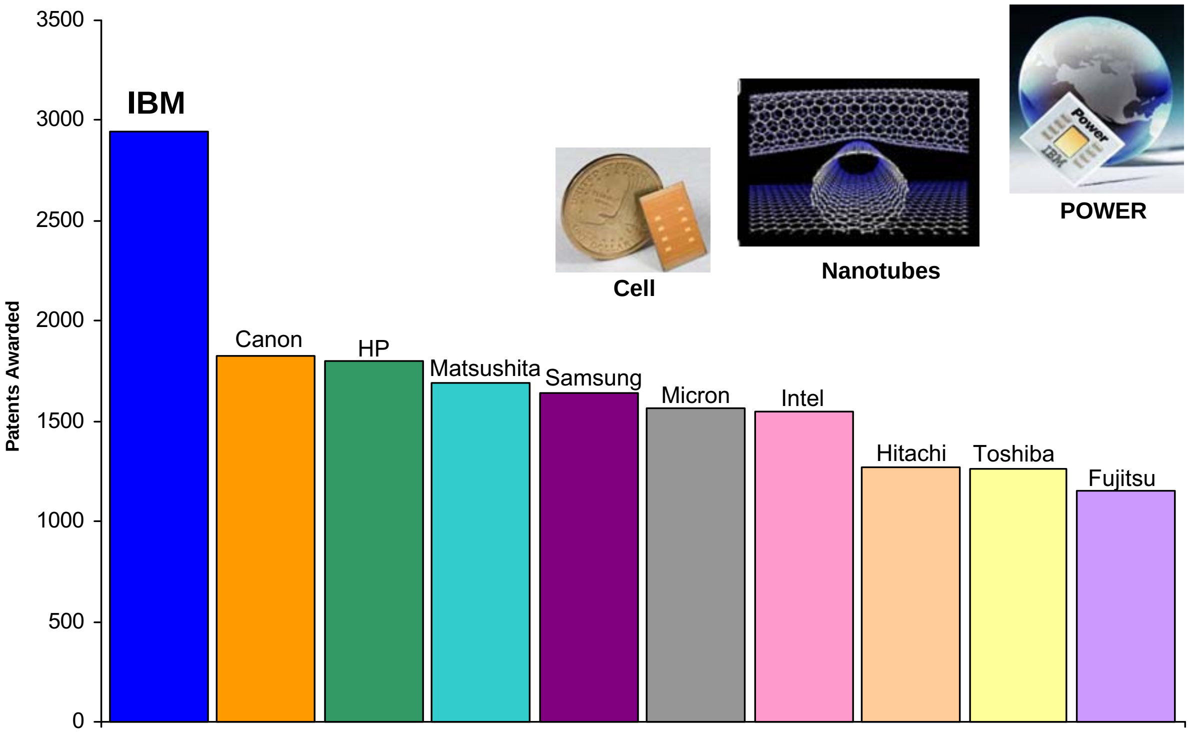

Figure 3-9. IBM's 2005 Patent Total: 13 Years of Leadership AM711.0

Notes:

IBM Power Architecture

From consumer electronics to supercomputers

A common architecture . . . the most scalable technology

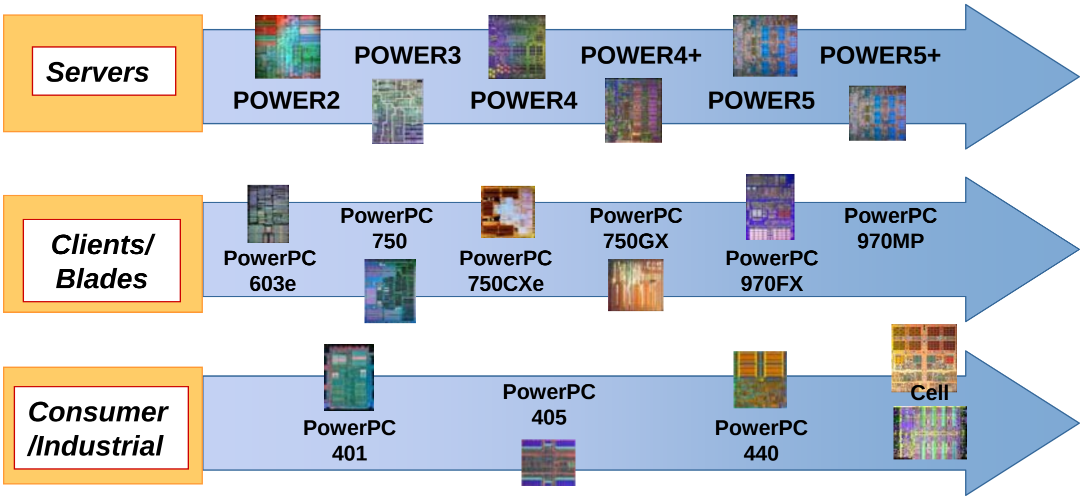

Servers

POWER2

POWER3

POWER4

POWER4+

POWER5

POWER5+

Clients/ Blades

PowerPC 603e

PowerPC 750

PowerPC 750CXe

PowerPC 750GX

PowerPC 970FX

PowerPC 970MP

Consumer /Industrial

PowerPC 401

PowerPC 405

PowerPC 440

Cell

Source: http://www.ibm.com/chips/power/aboutpower/

Figure 3-10. IBM Power Architecture AM711.0

Notes:

IBM is proud of the success of our POWER processor. Many industries have adopted it is extremely scalable technology... used in servers, desktop computing, and embedded in games…

From consumer electronics to supercomputers – the POWER architecture has been a winner from its first inception. And, you can rely on IBM continuing to provide technological advancements through the POWER technology…

IBM POWER Technology Roadmap for System p

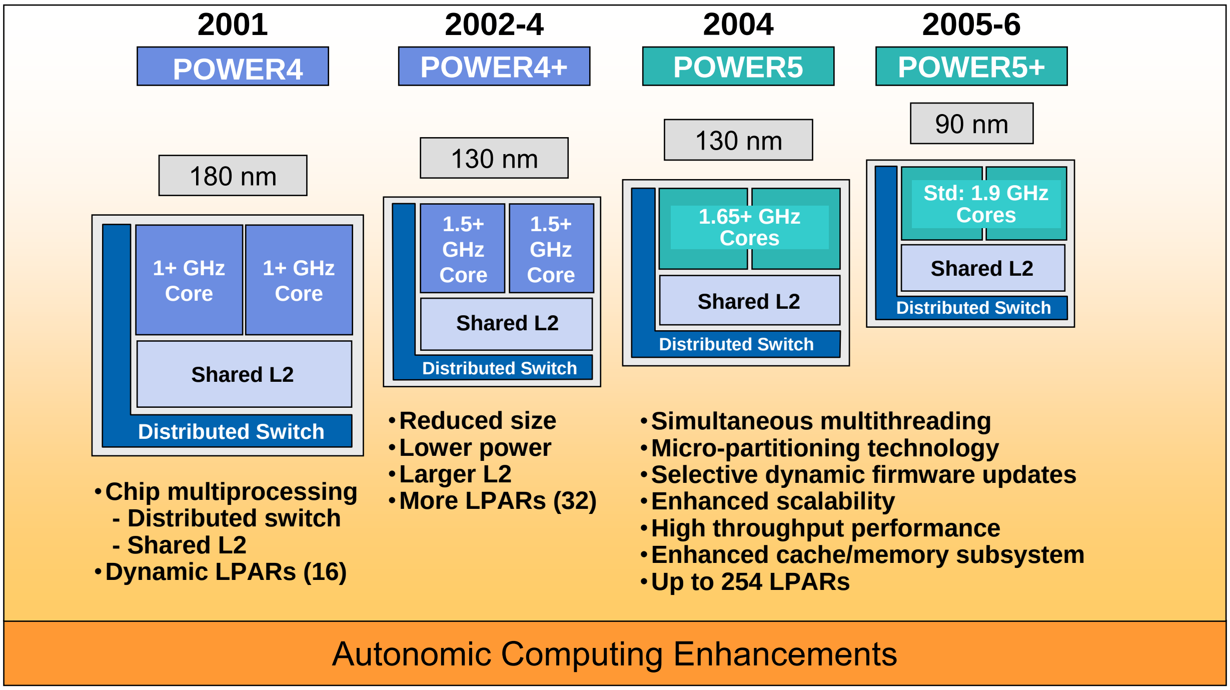

- Chip multiprocessing
  - Distributed switch
  - Shared L2
- Dynamic LPARs (16)

- Reduced size
- Lower power
- Larger L2
- More LPARs (32)

- Simultaneous multithreading
- Micro-partitioning technology
- Selective dynamic firmware updates
- Enhanced scalability
- High throughput performance
- Enhanced cache/memory subsystem
- Up to 254 LPARs

Autonomic Computing Enhancements

Figure 3-11. IBM POWER Technology Roadmap for System p

Notes:

POWER5+ Enhancements

- POWER5+ enhancements….
  - 15% higher frequencies and higher bus clock rates
  - 37% reduction in size
  - Reduction in power consumption
  - Large page size support
  - Memory controller improvements
  - Quad-core module support
  - Better performance

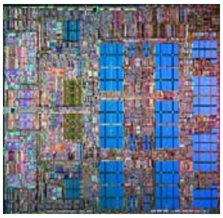

POWER5+

Figure 3-12. POWER5+ Enhancements AM711.0

Notes:

POWER5+ Quad-Core Module Highlights

Innovative Design
Leading Edge Implementation
High Efficiency
Excellent Price/Performance/Density

Double your performance per socket!

- First 4-core processor package in the industry
- Two dual-core POWER5+ chips on one substrate
- Multi-threaded CPU cores yield eight simultaneous threads per socket
- Featuring a massive 72 MB of high-speed cache (36 MB for each POWER5+ chip)
- In a module that is about the same physical size as other low-end modules
- Implemented using IBM's highly reliable MCM technology and multi-core CPUs

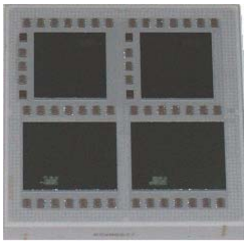

Quad-Core Module
QCM

Figure 3-13. POWER5+ Quad-Core Module Highlights AM711.0

Notes:

POWER5+ Packaging

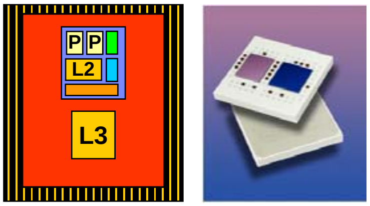

Dual-Core Module
POWER5+ dual-core chip + L3 cache chip
(520, 550)
Two processor cores

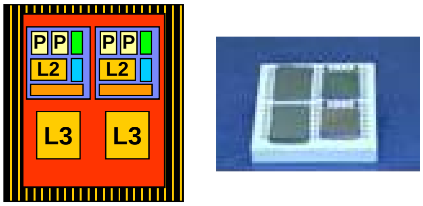

Quad-Core Module
Two POWER5+ dual-core chips + two L3 cache chips (550Q)
Four processor cores

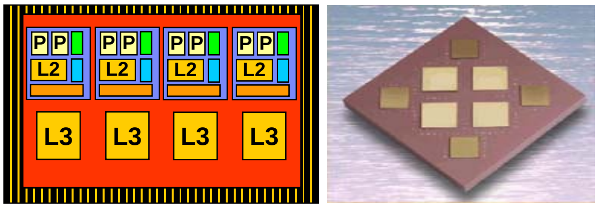

Multi-Chip Module
Four dual-core POWER5 chips + four L3 cache chips
(590, 595)
Eight processor cores

Figure 3-14. POWER5+ Packaging AM711.0

Notes:

POWER5 / 5+ Design

POWER5 design	**1.5, 1.65 and 1.9 GHz**	**.13 micron**
POWER5+ design	**1.5 and 1.9 GHz**	**.09 micron**

POWER5 / 5+

- Simultaneous multi-threading
- Micro-partitioning hardware support
  - Sub-processor allocation
- Enhanced distributed switch
- Enhanced memory subsystem
  - Large L3 cache: 36 MB
  - Memory controller on-chip
- Dynamic power saving
  - Clock gating

POWER5/5+

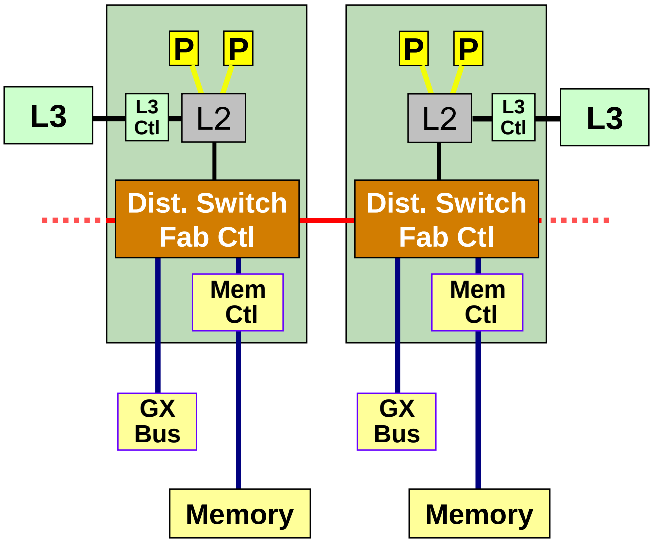

Figure 3-15. POWER5 / 5+ Design AM711.0

Notes:

IBM System p5: Simultaneous Multi-Threading

POWER5 (simultaneous multithreading)

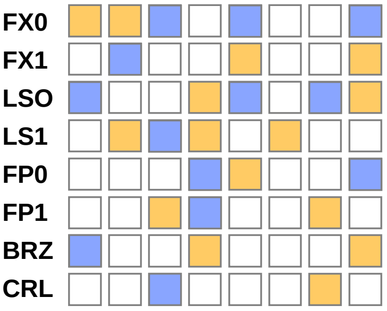

🟨 **Thread0 active**
☐ **No thread active**
🟦 **Thread1 active**

Appears as four CPUs per chip to the operating system (AIX 5L V5.3 and Linux)

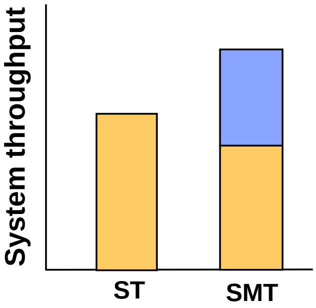

- Utilizes unused execution unit cycles
- Presents symmetric multiprocessing (SMP) programming model to software
- Natural fit with superscalar out-of-order execution core
- Dispatch two threads per processor: *"It is like doubling the number of processors."*
- Net result:
  - ***Better performance***
  - ***Better processor utilization***

Figure 3-16. IBM System p5: Simultaneous Multi-Threading

Notes:

POWER5 Performance

- Top 500 Supercomputers (November 2005)
  - IBM has positions 1, 2 and 3
  - IBM has 5 out of top ten
  - IBM badge on 219 entries 43.8% of Top 500
- TPC-C Benchmarks
  - Oracle 10g result on 32 way 595 beats HP 64 way
  - DB2 on 595 64 way almost double performance of HP 64 way

Figure 3-17. POWER5 Performance AM711.0

Notes:

http://www.top500.org/

The No. 1 position was again claimed by the BlueGene/L System, a joint development of IBM and DOE's National Nuclear Security Administration (NNSA) and installed at DOE's Lawrence Livermore National Laboratory in Livermore, California BlueGene/L also occupied the No. 1 position on the last TOP500 list issued in November 2004. However, the system was doubled in size during the last six months and reached a new record Linpack benchmark performance of 136.8 TFlop/s (“teraflops” or trillions of calculations per second). This system, once completed, will again be doubled in size and is expected to remain the #1 Supercomputer in the world for the next few editions of the TOP500 list. http://www.tpc.org/

Our POWER systems have delivered stunning benchmark results recently. Check the website for all the details.

Checkpoint Questions (1 of 7)

1. True or False. RISC architecture was invented by IBM.

2. The acronym RISC stands for:
   a. Reduced Instruction Set Computer
   b. Really Interesting System Concept
   c. Radically Intense Silly Computer
   d. None of the above

3. True or False. The acronym POWER stands for Performance Optimization with Enhanced RISC.

4. In measuring execution time of a program which of the following factors are important?
   a. Cycle time
   b. Compiler optimization
   c. Number of instructions/cycle
   d. All of the above are important factors

Figure 3-18. Checkpoint Questions (1 of 7) AM711.0

Notes:

Topic Summary

You should now be able to:

- Describe the advantages of RISC processor architecture.
- Explain the features of the POWER microprocessor family.
- Describe the processor roadmap for System p.
- Describe key enhancements in POWER5.

Figure 3-19. Topic Summary AM711.0

Notes:

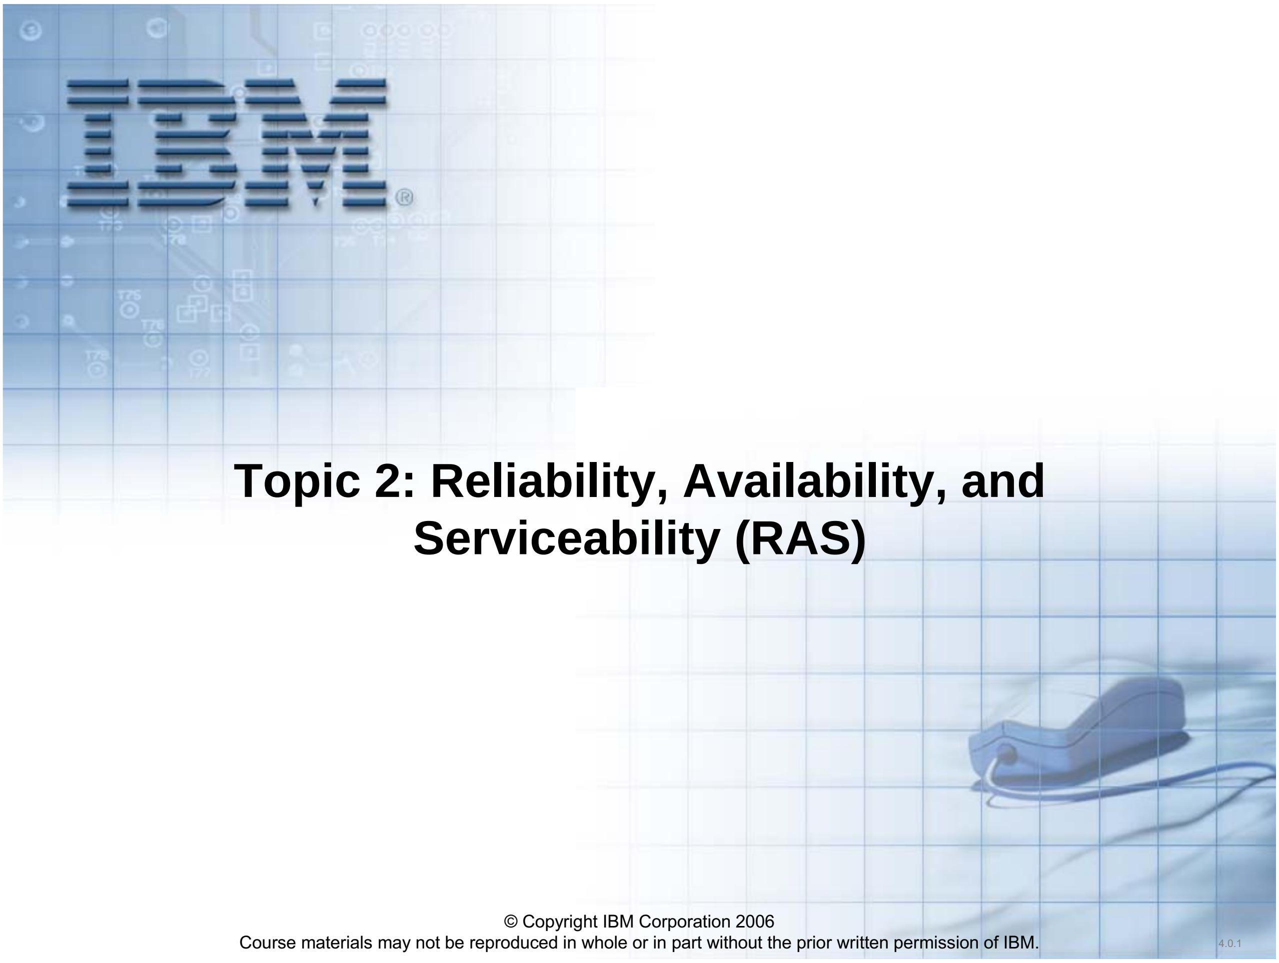

Topic 2: Reliability, Availability, and Serviceability (RAS)

Figure 3-20. Topic 2: Reliability, Availability, and Serviceability (RAS) AM711.0

Notes:

Topic Objectives

After completing this topic, you should be able to:

- Explain why RAS is important to our clients.
- Describe the autonomic computing initiative.
- List the features of System p that enhance RAS.

Figure 3-21. Topic Objectives AM711.0

Notes:

Why is RAS Important to Our Clients?

✓Increase return on investment

✓Improve quality of service

✓Improve time to value

© Copyright IBM Corporation 2006

Figure 3-22. Why is RAS Important to Our Clients? AM711.0

Notes:

Autonomic computing will bring significant return on investment. It will reduce the total IT spending such that customers can once again afford to invest in new technology. It will improve the utilization of IT staff, and will reduce “new technology” expenses by enabling more work to be done with current equipment. And it will enable our customers to align their IT organization tightly with their business goals.

Autonomic computing will bring improved quality of service by eliminating a significant percentage of system outages and security breaches. It will bring significantly improved resiliency and responsiveness to users.

Autonomic computing will also improve time to value since new solutions will be brought online much faster, thus making the business more responsive to changing market demands.

The new autonomic systems can be justified based on the return they provide to the business whether it is a reduction in internal cost, or increased revenue. And bringing them on-line faster simply means an earlier realization of these economic benefits.

What is the Autonomic Computing Initiative?

- Autonomic computing
  - One of the key components of an On Demand operating environment
  - Modelled on human autonomic nervous system

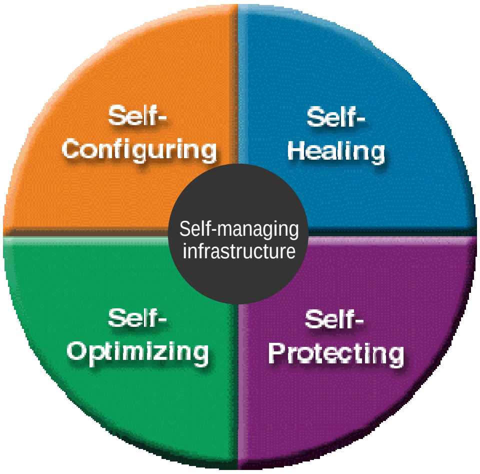

Figure 3-23. What is the Autonomic Computing Initiative? AM711.0

Notes:

Autonomic computing is a term coined by IBM to describe the need to shift the burden of managing IT systems from IT professionals to the systems themselves. The term comes from the autonomic nervous system of the human body. That is the system that regulates your body's basic functions without your conscious awareness.

For instance, when you need to run to catch a train, you don't need to consciously decide to excrete adrenaline, reallocate oxygen to the muscles in your legs, and increase your heart rate. Those important and necessary physical adjustments are handled for you automatically. In a similar way, IBM is driven to create IT systems that handle more and more tasks on their own, without the need for intervention on the part of the IT staff.

IBM has a strategy of delivering e-business “on demand.” On demand computing means enabling enterprises to have business processes -- integrated end-to-end across the company and with key partners, suppliers, and customers -- that can respond with speed to any customer demand, market opportunity, or external threat.

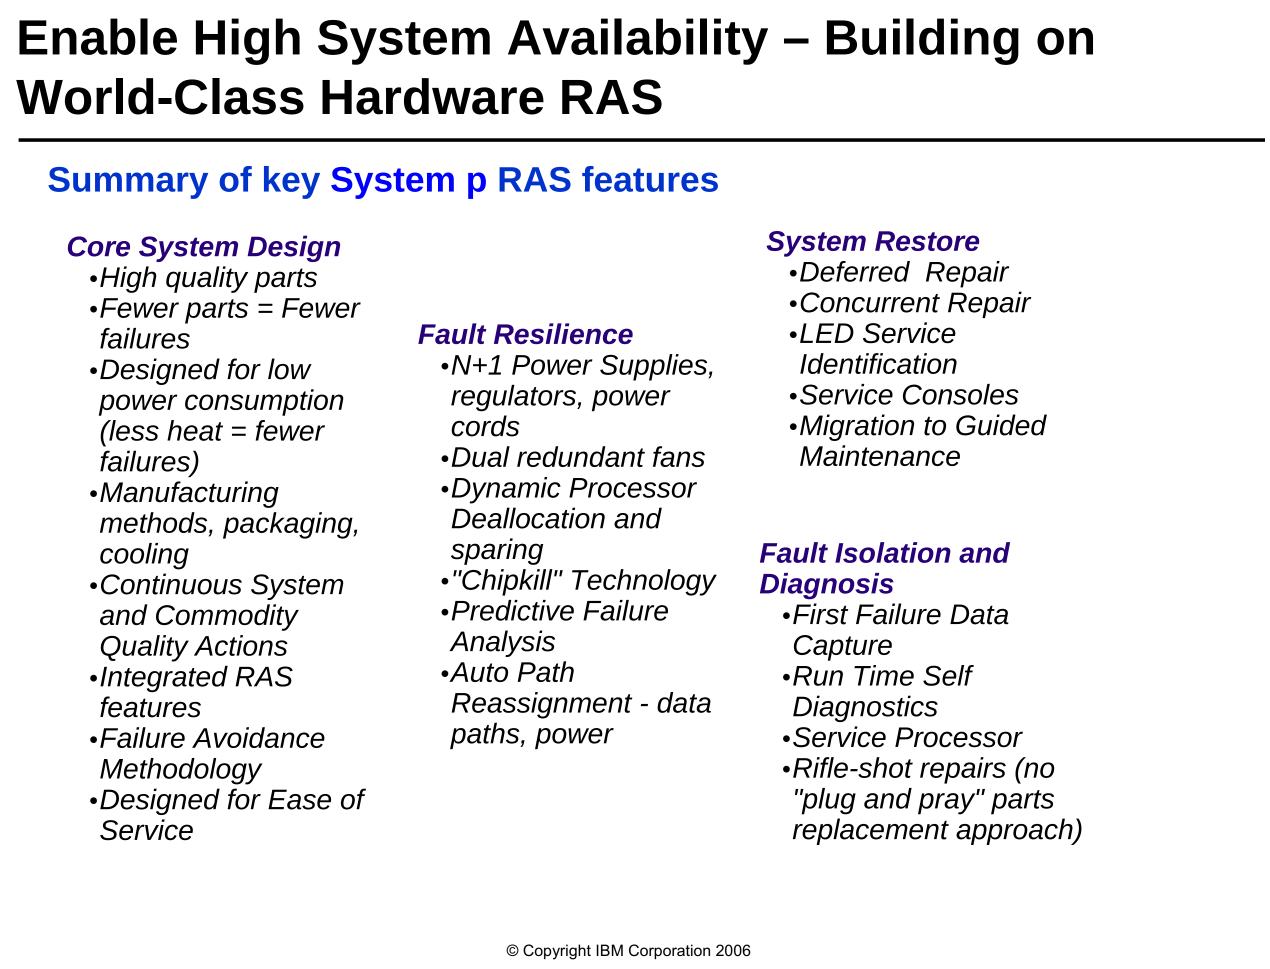

Enable High System Availability – Building on World-Class Hardware RAS

Summary of key System p RAS features

Core System Design

- High quality parts
- Fewer parts = Fewer failures
- Designed for low power consumption (less heat = fewer failures)
- Manufacturing methods, packaging, cooling
- Continuous System and Commodity Quality Actions
- Integrated RAS features
- Failure Avoidance Methodology
- Designed for Ease of Service

Fault Resilience

- N+1 Power Supplies, regulators, power cords
- Dual redundant fans
- Dynamic Processor Deallocation and sparing
- "Chipkill" Technology
- Predictive Failure Analysis
- Auto Path Reassignment - data paths, power

System Restore

- Deferred Repair
- Concurrent Repair
- LED Service Identification
- Service Consoles
- Migration to Guided Maintenance

Fault Isolation and Diagnosis

- First Failure Data Capture
- Run Time Self Diagnostics
- Service Processor
- Rifle-shot repairs (no "plug and pray" parts replacement approach)

Figure 3-24. Enable High System Availability – Building on World-Class Hardware RAS AM711.0

Notes:

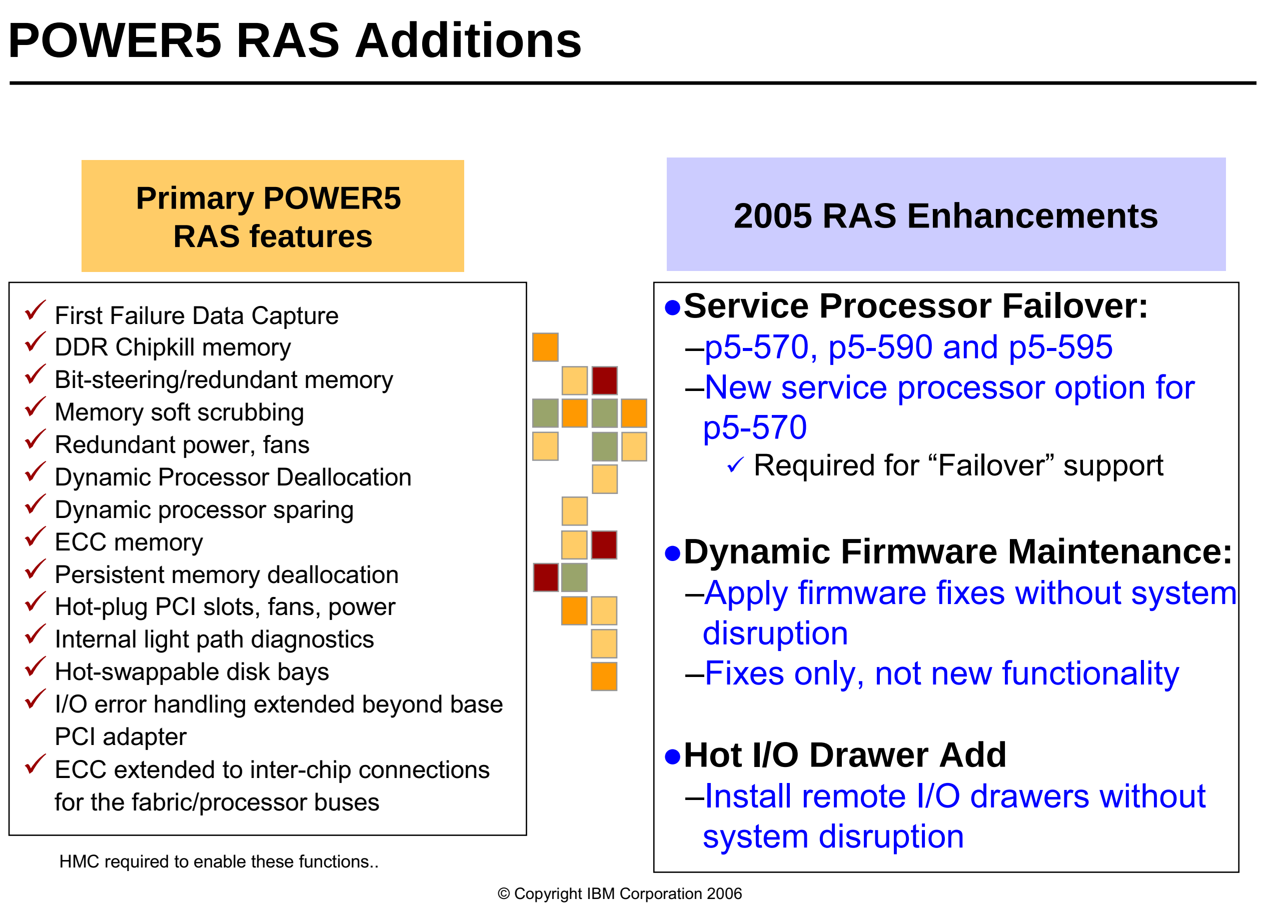

POWER5 RAS Additions

Primary POWER5 RAS features

- ✓ First Failure Data Capture
- ✓ DDR Chipkill memory
- ✓ Bit-steering/redundant memory
- ✓ Memory soft scrubbing
- ✓ Redundant power, fans
- ✓ Dynamic Processor Deallocation
- ✓ Dynamic processor sparing
- ✓ ECC memory
- ✓ Persistent memory deallocation
- ✓ Hot-plug PCI slots, fans, power
- ✓ Internal light path diagnostics
- ✓ Hot-swappable disk bays
- ✓ I/O error handling extended beyond base PCI adapter
- ✓ ECC extended to inter-chip connections for the fabric/processor buses

HMC required to enable these functions..

2005 RAS Enhancements

- **Service Processor Failover:**
  - –p5-570, p5-590 and p5-595
  - –New service processor option for p5-570
    - ✓ Required for “Failover” support
- **Dynamic Firmware Maintenance:**
  - –Apply firmware fixes without system disruption
  - –Fixes only, not new functionality
- **Hot I/O Drawer Add**
  - –Install remote I/O drawers without system disruption

Figure 3-25. POWER5 RAS Additions AM711.0

Notes:

Checkpoint Questions (2 of 7)

5. Which System p feature helps eliminate intermittent errors?
   a. Service processor
   b. Service agent
   c. First failure data capture
   d. Dynamic processor deallocation

6. Which of the following is not related to System p RAS?
   a. Light path diagnostics
   b. First failure data capture
   c. Simultaneous multi-threading
   d. Dynamic processor de-allocation

Figure 3-26. Checkpoint Questions (2 of 7)

Notes:

Checkpoint Questions (3 of 7)

7. Which of the following does not provide additional reliability for System p memory?
   a. Bit-scattering
   b. Bit-steering
   c. Bit-scrubbing
   d. Bit-of-this-bit-of-that

8. What is true about System p processors?
   a. They can be hot-swapped for convenience
   b. They never fail due to first failure data capture
   c. They may be implemented as hot spares
   d. They can be fixed via firmware

Figure 3-27. Checkpoint Questions (3 of 7) AM711.0

Notes:

Topic Summary

You should now be able to:

- Explain why RAS is important to our clients.
- Describe the autonomic computing initiative.
- List the features of System p that enhance RAS.

Figure 3-28. Topic Summary AM711.0

Notes:

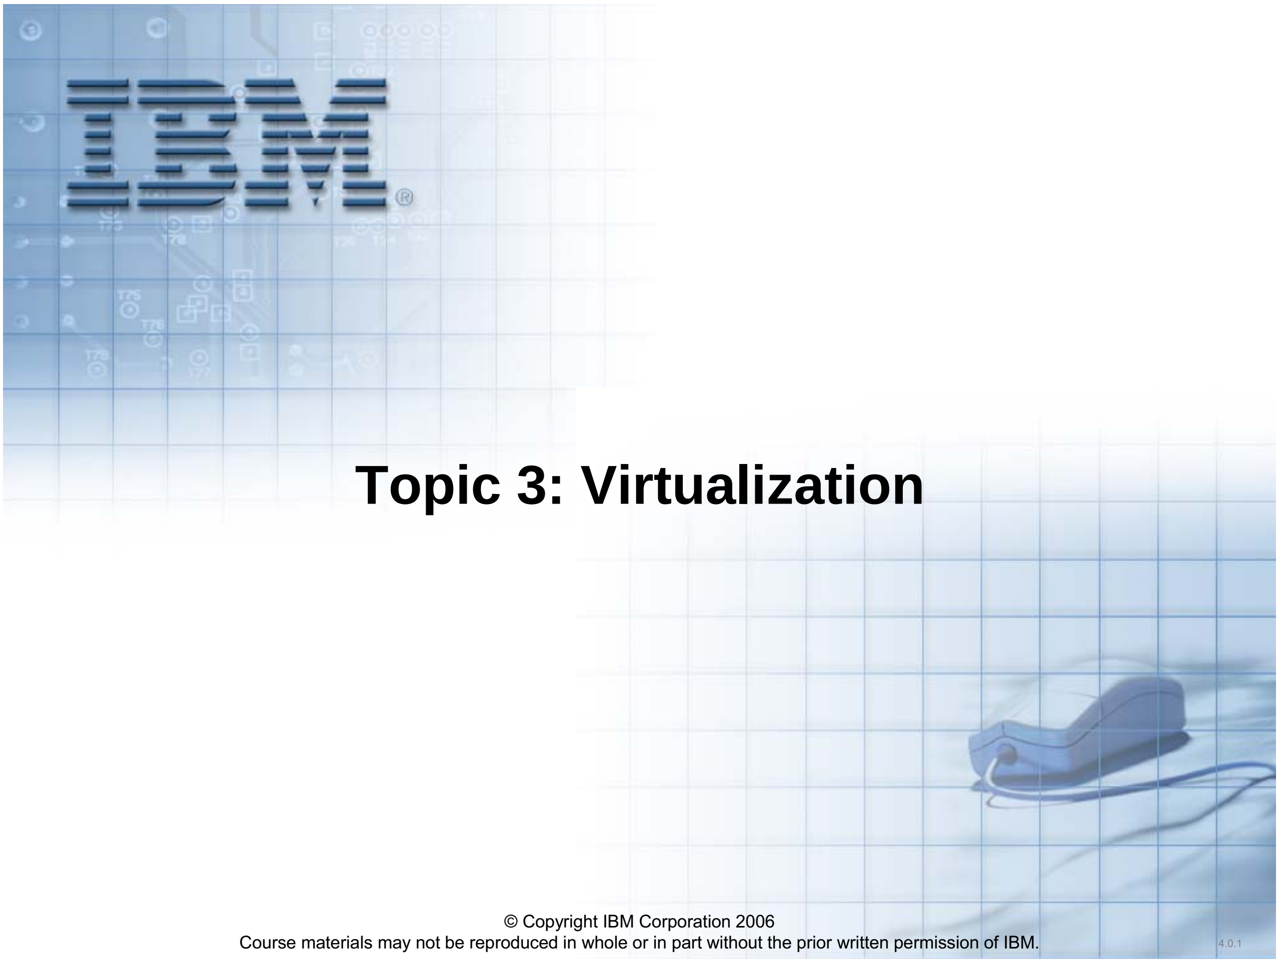

Topic 3: Virtualization

4.0.1

Figure 3-29. Topic 3: Virtualization AM711.0

Notes:

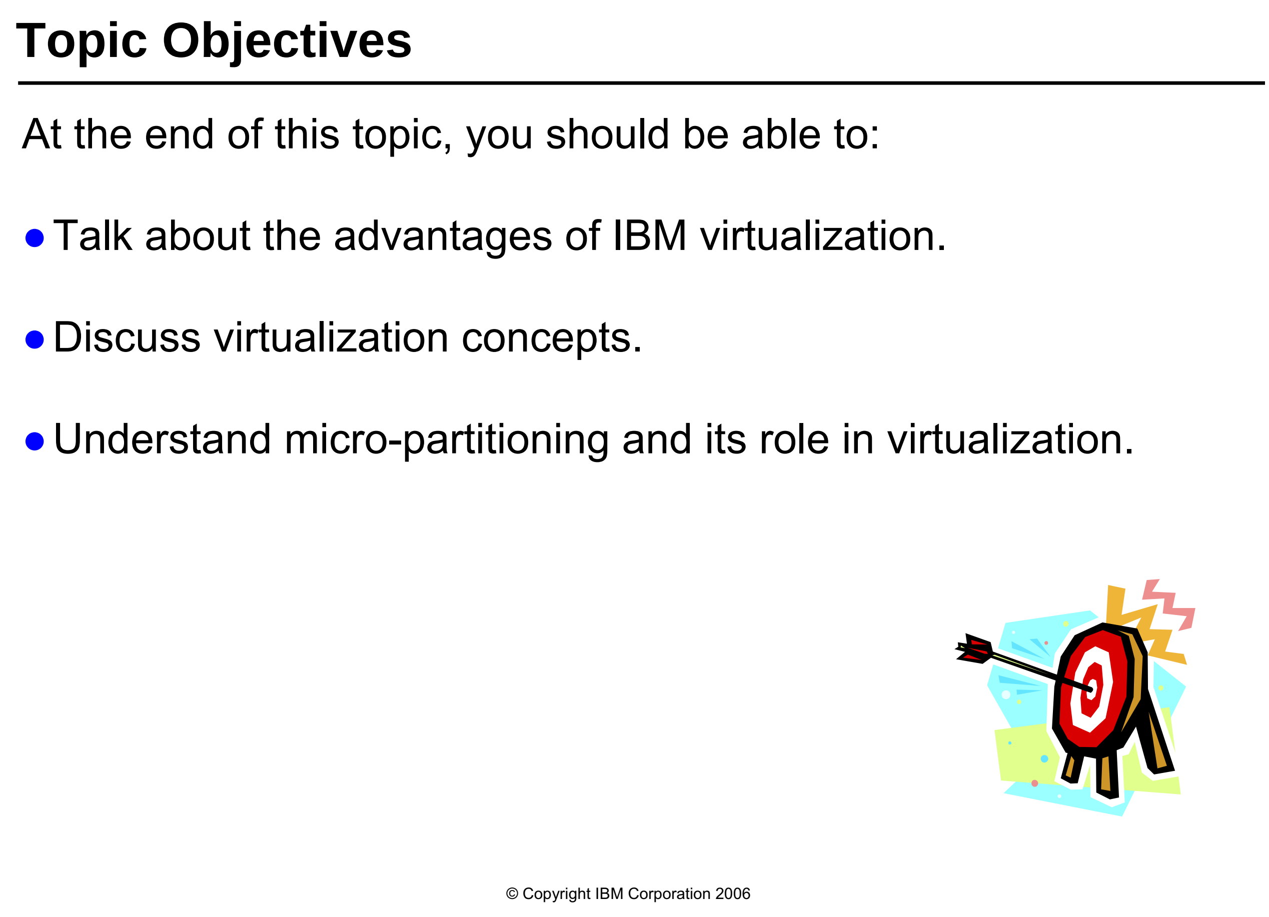

Topic Objectives

At the end of this topic, you should be able to:

- Talk about the advantages of IBM virtualization.
- Discuss virtualization concepts.
- Understand micro-partitioning and its role in virtualization.

Figure 3-30. Topic Objectives AM711.0

Notes:

Why IT Optimization and Virtualization?

1. **Complexity and administrative cost are growing faster than hardware cost.**
2. **Growing revenue while containing cost remains #1 concern of management.**
3. **Businesses cannot address the need for integration of people, and so forth, without first addressing the underlying infrastructure.**

Cost of People versus Spending on new systems

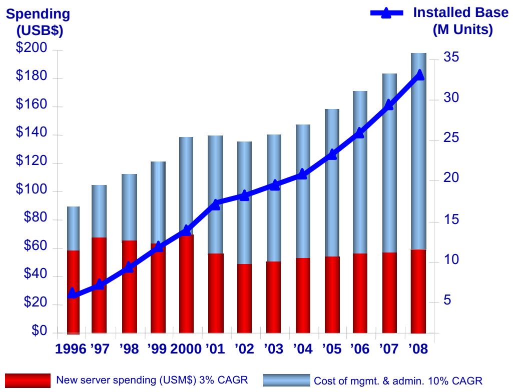

Source: IDC Directions 4-7-2004 Customer Adoption of On-Demand Enterprises

- ✓ Lower the cost of their existing infrastructure
- ✓ Reduce the complexity of adding to that infrastructure
- ✓ Build heterogeneous infrastructure across multiple Data Centers that are more responsive to their business needs

Figure 3-31. Why IT Optimization and Virtualization? AM711.0

Notes:

The purpose of this chart is to set the business concerns and how Virtualization solutions address the concern. In the chart on the right, the red reflects spending on new systems, which you can see over time has leveled off and is even down somewhat.

The blue bars of the chart reflect the cost of people, which keeps escalating, and is expected to more than quadruple by 2008. The blue trend line in the graph shows a rapid growth of server footprints and the associated management expense. Virtualization can help reduce the number of footprints by increasing the utilization of servers and reallocating resources.

Increasing utilization of information assets, simplifying the IT infrastructure and reducing operating costs across servers, storage, networking, and grid computing is what IBM Virtualization Solutions is all about. By extracting some of that administrative cost out the infrastructure and by increasing system and resource utilization and improving productivity and flexibility, these "virtualized" IT assets can help fuel the business growth we just talked about, control the cost in doing so, and increase staff productivity at the same time.

What is System Utilization and Efficiency?

Answer: Ability to maintain workload responsiveness at the lowest system cost.

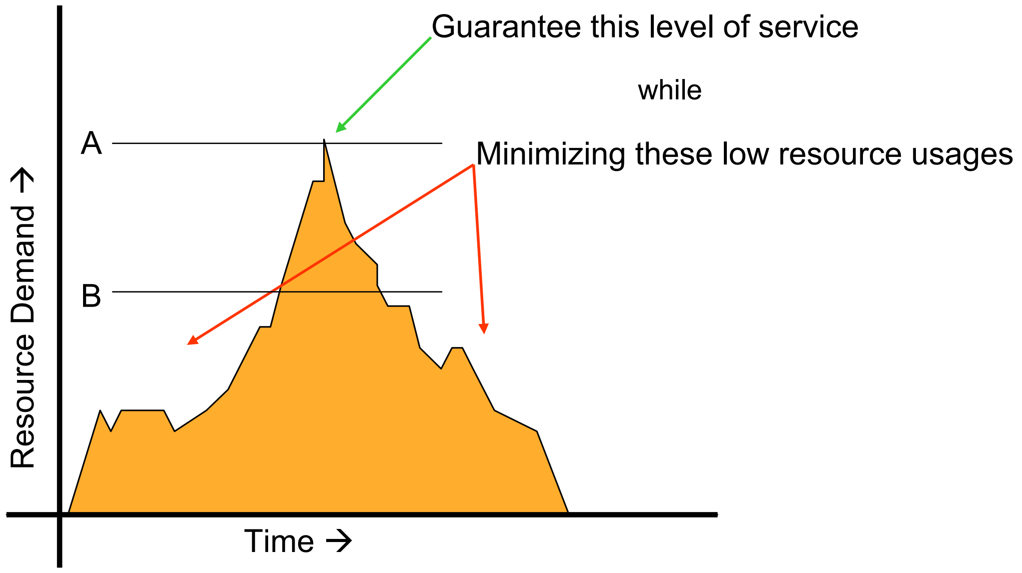

Guarantee this level of service

while

Minimizing these low resource usages

Figure 3-32. What is System Utilization and Efficiency? AM711.0

Notes:

Virtualization

"Virtualization is the process of presenting computing resources in ways that users and applications can easily get value out of them, rather than presenting them in a way dictated by their implementation, geographic location, or physical packaging. In other words, it provides a logical rather than physical view of data, computing power, storage capacity, and other resources."

Jonathan Eunice, Illuminata Inc.

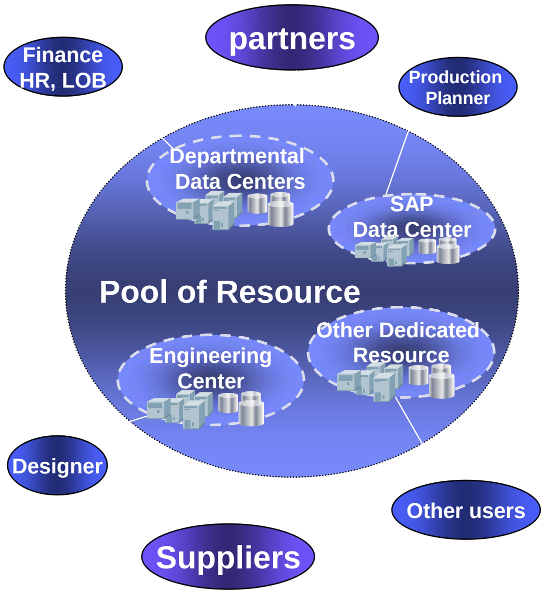

Figure 3-33. Virtualization AM711.0

Notes:

While Virtualization sounds like a very complex, technical thought, it is really a simple idea.

In IT, as Jonathan Eunice from Illuminata, puts it, we do a similar thing when we take a logical rather than a physical view of your resources. For example, in storage, rather than saying we have three different types of storage systems, which together might total 30 Tb of disk space – we start managing the 430 Tb as a single type of resource. We focus on how to use the resource not how to manage it.

It's a technique we have been using in large mainframe computer for 30+ years. Not having to manage each computer or resource separately – but to mange them together, virtually, allows for huge improvements in utilization. A typical mainframe today runs at between 70% and 90% utilized. The rest of a companies infrastructure is probably running at less than 15% utilized. Rising the level of utilization across your whole infrastructure usually means you will need to manage less things. Less things require less people.

So, and increase staff productivity at the same time. Having simplified its then easy to automate which can reduce errors and create streamlined-business responsive systems.

Value of Virtualization

IBM Virtualization Benefits

- ☑ *Can help improve cost and speed*
- ☑ *Allows easier provisioning, management and deployment*
- ☑ *Changes are transparent to users, which helps improve efficiency*

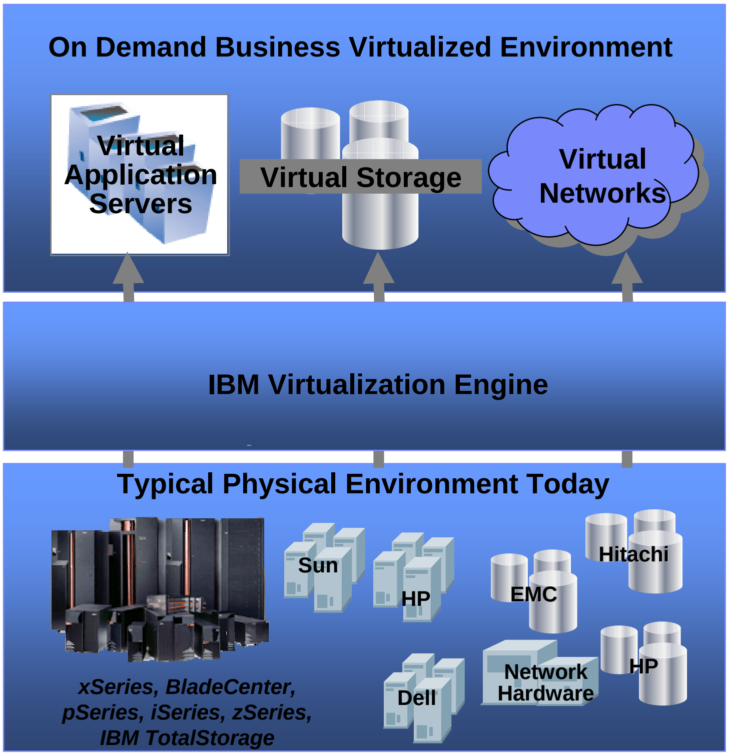
On Demand Business Virtualized Environment
Virtual Application Servers
Virtual Storage
Virtual Networks
IBM Virtualization Engine
Typical Physical Environment Today
xSeries, BladeCenter, pSeries, iSeries, zSeries, IBM TotalStorage
Sun
HP
Dell
EMC
Hitachi
Network Hardware
HP

Figure 3-34. Value of Virtualization AM711.0

Notes:

IBM Virtualization Engine is the Technology and Services Framework to Help Clients Implement Virtualization

Key Components of the Virtualization Engine suite for Servers:

Tivoli Provisioning Manager

- Adds, deletes, moves and configures servers, software and network resources
- Automates resource setup and configuration
- Executes IT processes in a consistent, customized and error-free manner

IBM Director

- Provides a simplified view of IBM System p hardware from a centralized server
- Predictive, proactive capabilities for real-time system diagnostics

IBM Enterprise Workload Manager

- Automatic monitoring and management of heterogeneous workloads
- Identify and track work requests based on service class definitions

IBM Grid Toolbox

- Toolkit for creating and hosting grid services
- Supports Open Grid Services Architecture

Virtualization Engine Console

- Consolidates the view of enterprise resources
- Acts as a launch point for VE services

Figure 3-35. IBM Virtualization Engine is the Technology and Services Framework to Help Clients Implement Virtualization AM711.0

Notes:

Advanced POWER Virtualization Option

- Virtual I/O Server
  - Shared Ethernet
  - Shared SCSI and Fibre Channel-attached disk subsystems
  - Supports AIX 5L V5.3 and Linux* partitions
  - Managed via HMC or IVM
- Micro-Partitioning
  - Share processors across multiple partitions
  - Minimum partition 1/10th processor
  - AIX 5L V5.3, Linux*, or i5/OS**
- Partition Load Manager
  - Both AIX 5L V5.2 and AIX 5L V5.3 supported
  - Balances processor and memory request

*** SLES 9 or RHEL AS 3 and above**

**Available on 1.65 GHz p5-570, p5-590 and p5-595 models

Figure 3-36. Advanced POWER Virtualization Option AM711.0

Notes:

LPAR and Dynamic LPAR

- Standard on all new System p systems
- Whole processor allocations in LPARs

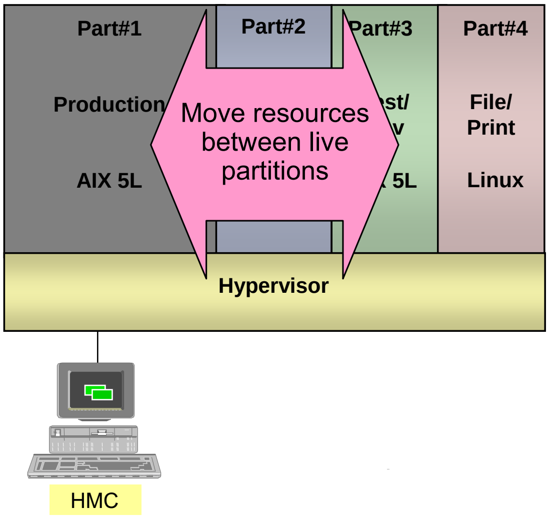

Figure 3-37. LPAR and Dynamic LPAR AM711.0

Notes:

Allocate processors, memory and I/O to create virtual servers

Minimum 128 MB memory, one CPU

All resources can be allocated independently

Resources can be moved between live partitions

Applications notified of configuration changes

Movement can be automated using Partition Load Manager

Works with AIX or Linux or iOS (590 and 595 only and max. of two processors)

Shared Processor Partitions

- Micro-Partitioning allows for multiple partitions to share one physical processor
- Up to 10 partitions per physical processor
- Up to 254 partitions active at the same time
  - Maximum per server
  - Maximum per HMC
- Partition's resource definition
  - Minimum, desired, and maximum values for each resource
  - Processor capacity
  - Virtual processors
  - Capped or uncapped
    - Capacity weight
  - Dedicated memory
    - Minimum of 128 MB and 16 MB increments
  - Physical or virtual I/O resources

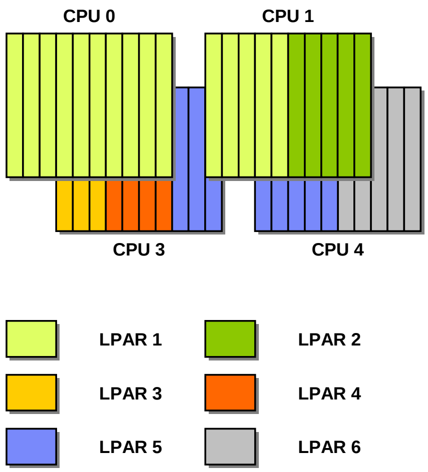

Figure 3-38. Shared Processor Partitions AM711.0

Notes:

Micro-partitioning allows for multiple partitions to share one physical processor.

A partition may be defined with a processor capacity as small as .1 processor units. This represents 1/10 of a physical processor. Each processor can be shared by up to 10 shared processor partitions. The shared processor partitions are dispatched and time-sliced on the physical processors under control of the POWER Hypervisor.

Micro-partitioning is supported across the entire POWER5 product line from the entry to the high-end systems.

Shared processor partitions still need dedicated memory, but the partitions I/O requirements can be supported through Virtual Ethernet and Virtual SCSI Server. Utilizing all virtualization features support for up to 254 shared processor partitions is possible.

The shared processor partitions are created and managed by the HMC. When you start creating a partition, you have to choose between a shared processor partition and a dedicated processor partition.

Capped and Uncapped Partitions

- Capped partition
  - Not allowed to exceed its capacity entitlement (CE)
- Uncapped partition
  - Is allowed to exceed its entitlement
- Capacity weight
  - Used for prioritizing uncapped partitions
  - Value 0-255
  - Value of 0 referred to as a “soft cap”

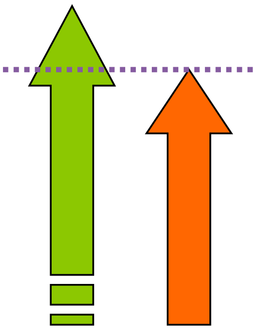

Figure 3-39. Capped and Uncapped Partitions AM711.0

Notes:

A capped partition is not allowed to exceed its capacity entitlement, while an uncapped partition is. In fact, it may exceed its maximum processor capacity. An uncapped partition is only limited in its ability to consume cycles by the lack of online virtual processors and its variable capacity weight attribute.

The variable capacity weight attribute is a number between 0–255, which represents the relative share of extra capacity that the partition is eligible to receive. This parameter applies only to uncapped partitions. A partition's share is computed by dividing its variable capacity weight by the sum of the variable capacity weights for all uncapped partitions. Therefore, a value of 0 may be used to prevent a partition from receiving extra capacity. This is sometimes referred to as a “soft cap.”

Shared Processor LPAR - Responsive Partition Example

- Work can be **prioritized** to give **time critical tasks** a high fraction of resource
- **Lower priority partitions** can run with **less dedicated resource**, but can “soak” up spare CPU cycles **automatically**

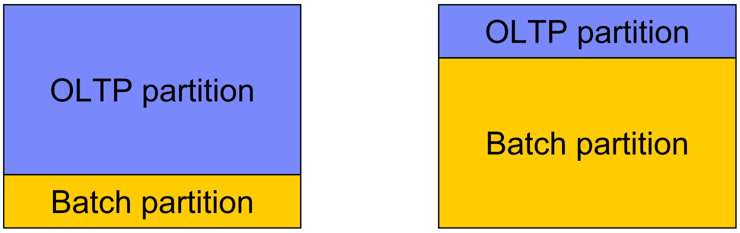

High interactive load gets precedence when needed

When interactive load is **low**, batch partition uses available resources

Figure 3-40. Shared Processor LPAR - Responsive Partition Example AM711.0

Notes:

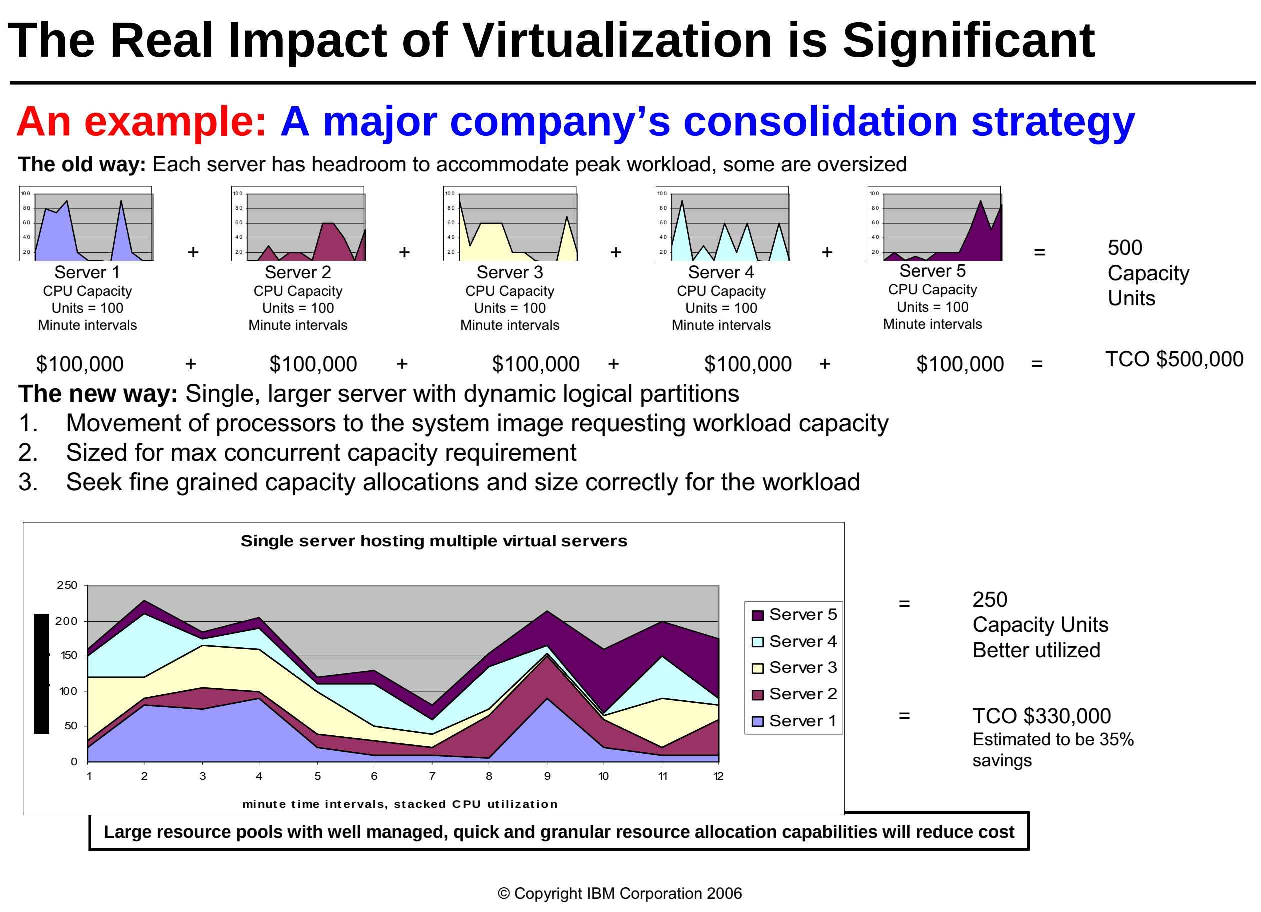

The Real Impact of Virtualization is Significant

An example: A major company’s consolidation strategy

The old way: Each server has headroom to accommodate peak workload, some are oversized

Server 1 CPU Capacity Units = 100 Minute intervals + Server 2 CPU Capacity Units = 100 Minute intervals + Server 3 CPU Capacity Units = 100 Minute intervals + Server 4 CPU Capacity Units = 100 Minute intervals + Server 5 CPU Capacity Units = 100 Minute intervals = 500 Capacity Units

$100,000 + $100,000 + $100,000 + $100,000 + $100,000 = TCO $500,000

The new way: Single, larger server with dynamic logical partitions

1. Movement of processors to the system image requesting workload capacity
2. Sized for max concurrent capacity requirement
3. Seek fine grained capacity allocations and size correctly for the workload

100
80
60
40
20
Single server hosting multiple virtual servers
250
200
150
100
50
0
1
2
3
4
5
6
7
8
9
10
11
12
minute time intervals, stacked CPU utilization
Server 5
Server 4
Server 3
Server 2
Server 1

= 250 Capacity Units Better utilized

= TCO $330,000 Estimated to be 35% savings

Large resource pools with well managed, quick and granular resource allocation capabilities will reduce cost

Figure 3-41. The Real Impact of Virtualization is Significant AM711.0

Notes:

Virtual I/O

- Allows sharing of network and storage devices
- Vital for shared processor partitions
  - Overcomes potential limit of adapter slots due to high number of possible Micro-Partitions
  - Allows the creation of logical partitions without the need for additional physical resources
- Allows attachment of previously unsupported solutions in selected OS clients (for example, Linux)

Figure 3-42. Virtual I/O AM711.0

Notes:

The virtualization features of the POWER5 platform support up to 254 partitions, while the biggest planned server only provides up to 160 I/O slots. With each partition requiring at least one I/O slot for disk attachment and another one for network attachment, this puts a constraint on the number of partitions. To overcome these physical limitations, I/O resources have to be virtualized. Virtual SCSI provides the means to do this for storage devices.

On the other hand, virtual I/O has a value proposition to it. It allows the creation of logical partitions without the need for additional physical resources. This facilitates on demand computing and server consolidation. Virtual I/O also provides a more economic I/O model by using physical resources more efficiently through sharing.

Furthermore, virtual I/O allows attachment of previously unsupported storage solutions. As long as the virtual I/O server supports the attachment of a storage resource, any client partition can access this storage by using virtual SCSI adapters.

Virtualization Simplifies Infrastructure

No Virtualization **Virtualization**

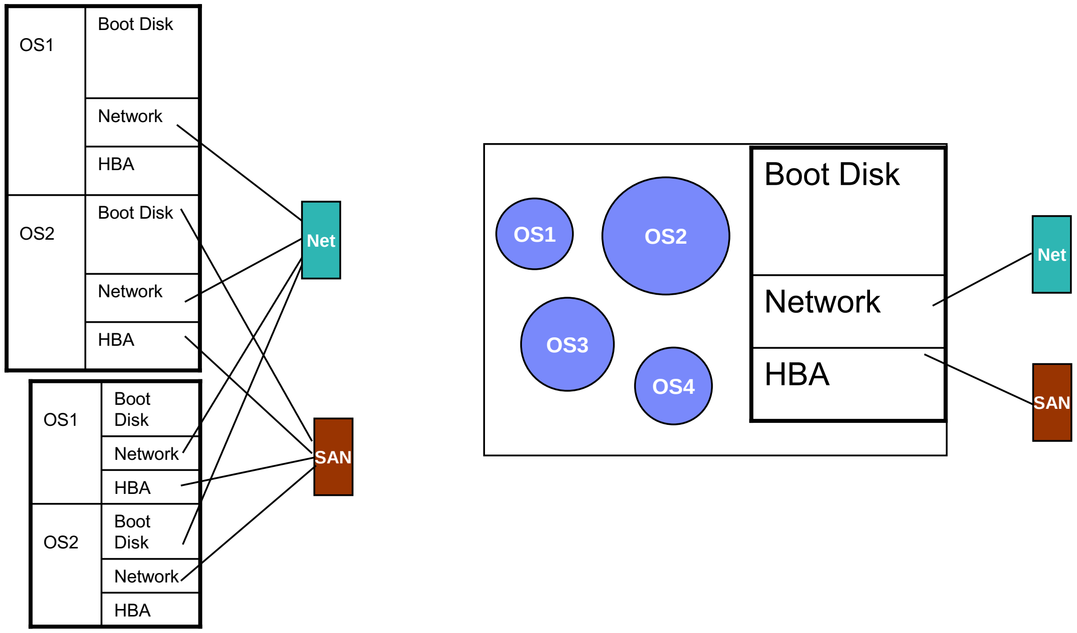

Dedicated Resources **Shared Resources**

Figure 3-43. Virtualization Simplifies Infrastructure AM711.0

Notes:

Virtual Networking

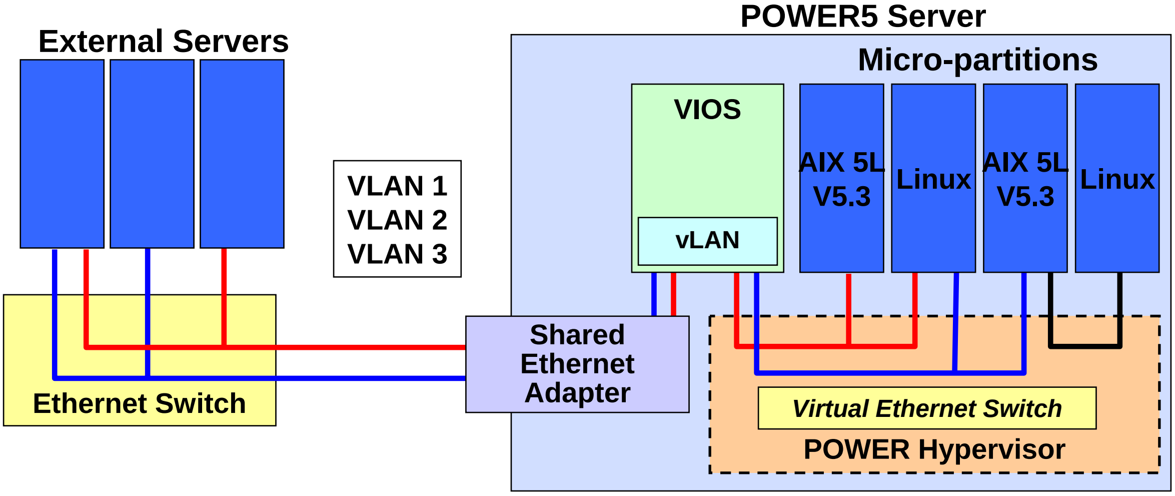

External Servers

VLAN 1
VLAN 2
VLAN 3

Ethernet Switch

POWER5 Server

Micro-partitions

VIOS

vLAN

AIX 5L V5.3

Linux

AIX 5L V5.3

Linux

Shared Ethernet Adapter

Virtual Ethernet Switch

POWER Hypervisor

Virtual Ethernet helps reduce hardware costs by eliminating LAN adapters

Shared Ethernet helps further reduce hardware costs by sharing LAN adapter resources

Available via optional Advanced POWER Virtualization or POWER Hypervisor and VIOS features.

Figure 3-44. Virtual Networking AM711.0

Notes:

Virtual Ethernet

- Partition to partition communication
- Requires AIX 5L V5.3 and POWER5

Shared Ethernet Adapter

- Provides access to outside world
- Uses Physical Adapter in the Virtual I/O Server

VLAN – Virtual LAN

- Provides ability for one adapter to be on multiple subnets
- Provides isolation of communication to VLAN members
- Allows a single adapter to support multiple subnets

Virtual SCSI

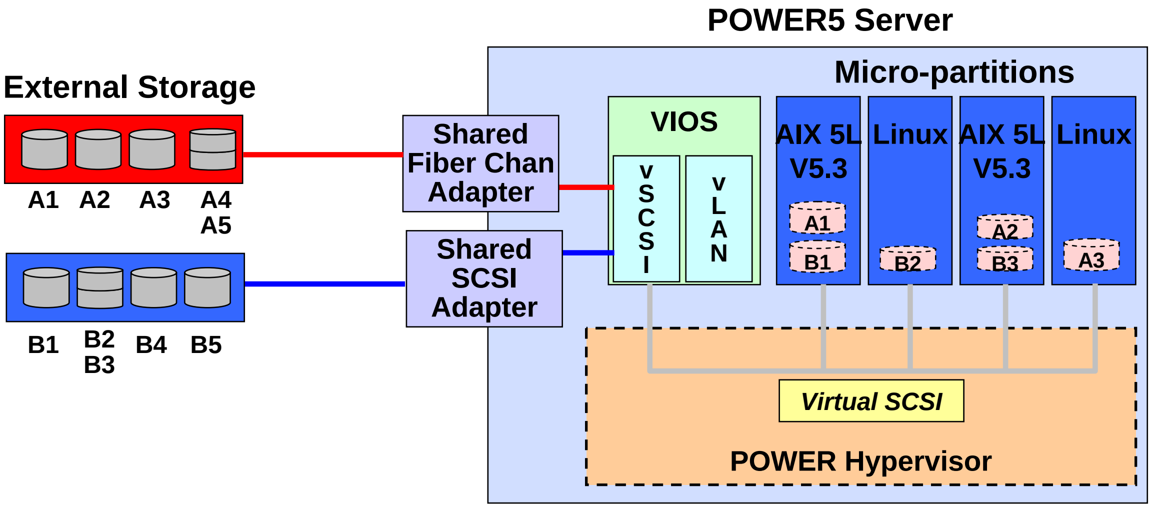

Virtual I/O helps reduce hardware costs by sharing disk drives

Available via optional Advanced POWER Virtualization or POWER Hypervisor and VIOS features.

Figure 3-45. Virtual SCSI AM711.0

Notes:

- VIOS owns physical disk resources
  - LVM based storage on VIO Server
  - Physical Storage can be SCSI or FC
    - Local or remote
- Micro-partition sees disks as vSCSI (Virtual SCSI) devices
  - Virtual SCSI devices added to partition via HMC
  - LUNs on VIOS accessed as vSCSI disk
  - VIOS must be active for client to boot
- Multiple LPARs can use same or different physical disk
  - Configure as logical volume on VIOS
  - Appear as a hdisk on the micro-partition
  - Can assign entire hdisk to a single client

Partition Load Manager

- Policy-based, automatic partition resource tuning
- Dynamically adjust CPU, memory allocation
- Supported on AIX 5L V5.3 and V5.21
- Supported on POWER4 and POWER5 systems

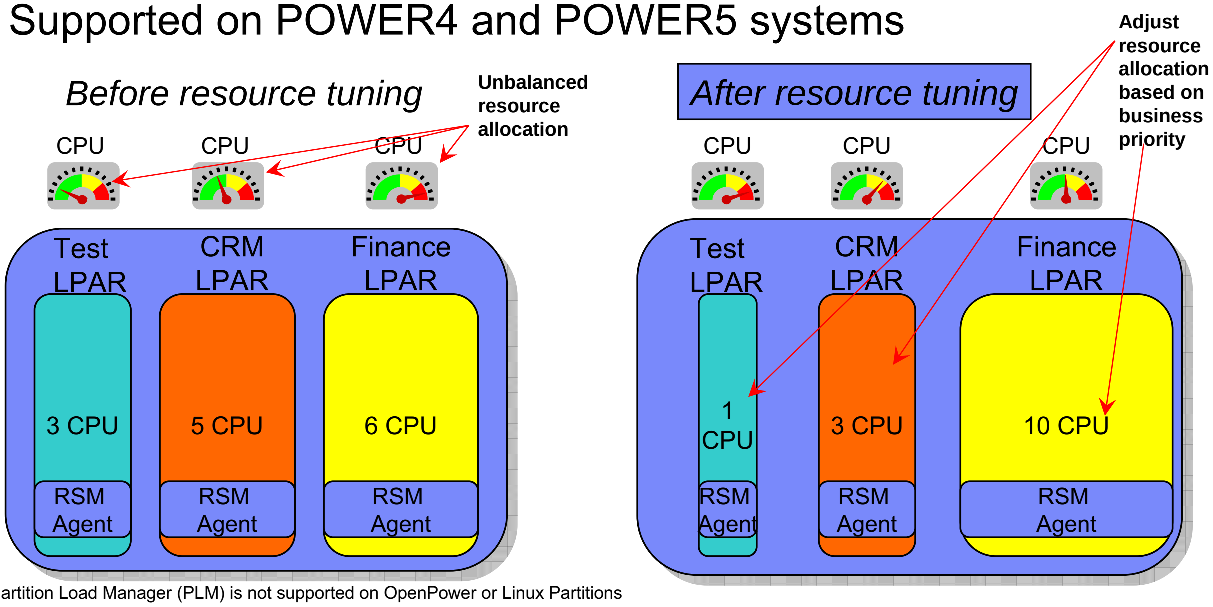

1 Partition Load Manager (PLM) is not supported on OpenPower or Linux Partitions

Figure 3-46. Partition Load Manager AM711.0

Notes:

Workload manager now supports automatically switching resource polices based on time of day - reducing administrator workload and helping to maximize system utilization.

Workload manager can enforce per process CPU, I/O and connect time limits by warning and killing processes that exceed these limits- protecting the system against errant or malicious processes. Also, the administrator can set up per class limits on the number of threads, processes or logins. The administrator can be automatically notified if these thresholds are reached.

“Processor Sets” allow the administrator to specify which physical processors are assigned to a process or class, Memory affinity provides the same function for memory - insuring maximum performance for high performance computing workloads that need to minimize memory latency. Note that processor sets are only applicable to SMP environments, not LPAR. Additional WLM resources were added to broaden the capabilities to manage diverse workloads.

Integrated Virtualization Manager (IVM)

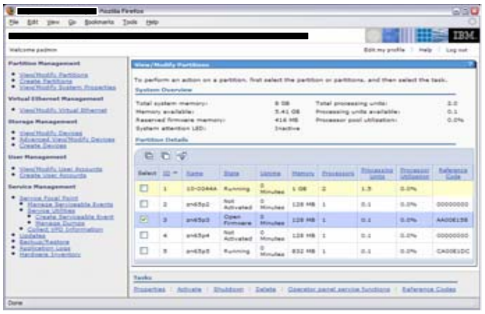

"HMC within a partition"

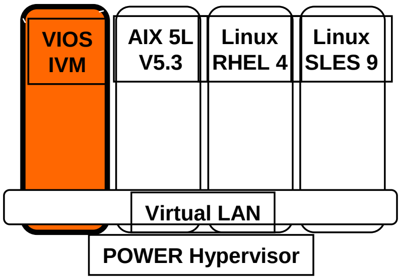

VIOS
IVM

AIX 5L
V5.3

Linux
RHEL 4

Linux
SLES 9

Virtual LAN

POWER Hypervisor

- Provides LPAR/virtualization support ***without a physical HMC*** – no HMC hardware translates to **lower entry point**
- **Web-based**, intuitive/**user-friendly** interface
- Shipped with the Virtual I/O Server (VIOS) – provides **tight integration** of LPAR management and virtual storage configuration
- Supports **creating/management of I/O and LPARs** within a single physical server
- All I/O is virtualized – Virtual Console, Storage, Ethernet, and Optical
- Subset of HMC Service functionality
  - •For example, no automatic DLPAR
  - •For example, single VIOS partition
  - •For example, maximum four VLANs
- Available for entry servers **up to p5-560Q Express** (both System p5 and @server p5 families)

Figure 3-47. Integrated Virtualization Manager (IVM) AM711.0

Notes:

5.2 is not supported in a virtual environment

OS Support for Virtualization

Function	AIX 5.2F	AIX 5.3	Linux SLES 9	Linux RHEL AS4
Dynamic LPAR				
• Processor	✓	✓	✓	✓
• Memory	✓	✓	✗	✗
• I/O	✓	✓	✓	✓
Virtual LAN	✗	✓	✓	✓
Micro-partitions	✗	✓	✓	✓
Virtual Storage	✗	✓	✓	✓
Virtual Ethernet	✗	✓	✓	✓
Partition Load Manager	✓	✓	✗	✗

i5/OS support available for 1.65 GHz p5-570, p5-590, p5-595

- i5/OS can participate in the shared processor pool, but only to the processor maximum
- I5/OS can utilize VLAN, but can NOT be a client to VIO server

NOTES:

- Dynamic LPAR standard on System p5, @server p5; part of POWER Hypervisor on OpenPower
- Virtual I/O Hosting Partition is an “appliance” – encapsulated partition available on System p5, @server p5 and OpenPower

Figure 3-48. OS Support for Virtualization AM711.0

Notes:

AIX 5.3 provides support for all the new features.

Linux support will continue to improve as the Linux kernel gets updated to provide the under-lying enablement

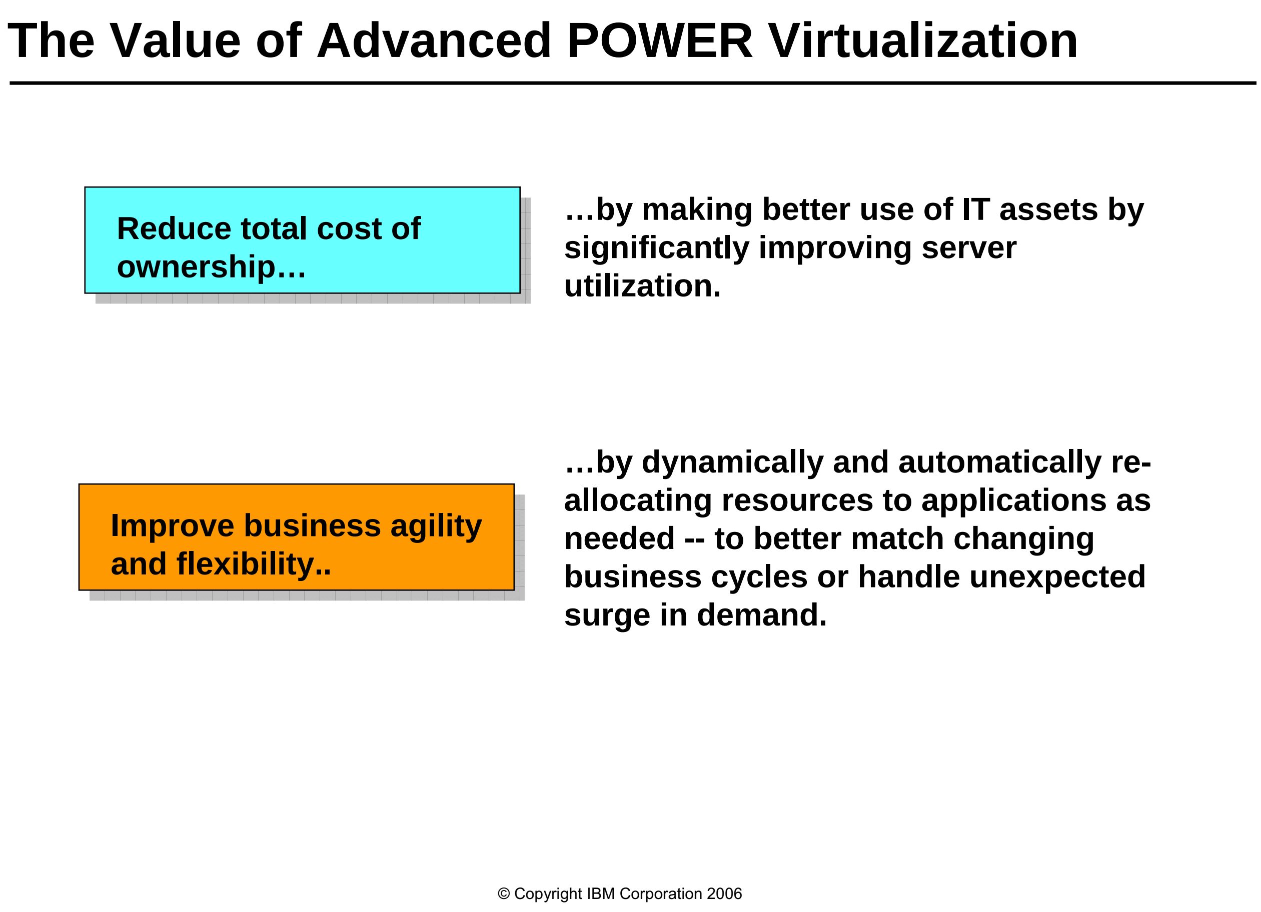

The Value of Advanced POWER Virtualization

Reduce total cost of ownership…

…by making better use of IT assets by significantly improving server utilization.

Improve business agility and flexibility..

…by dynamically and automatically re-allocating resources to applications as needed -- to better match changing business cycles or handle unexpected surge in demand.

Figure 3-49. The Value of Advanced POWER Virtualization AM711.0

Notes:

Consolidation of under-utilized (for example, 10-15%) workloads (onto partitions with <1 CPU)

Approach mainframe inspired utilization without sacrificing resource availability

Economic I/O model (for example, Shared Ethernet, Shared storage)

Simplify IT management (for example, rapid deployment of additional capacity)

Support multiple workloads/departments/applications

Rapid response (10 ms) to changing workload needs via shared processor pool

Automatic access to spare capacity via APV/Reserve CoD

Exploitation of VLAN (for example, high-speed communication)

Scenarios – Where APV is Likely to Have Significant Impact

NOT EXHAUSTIVE

Server consolidation of many, small underutilized servers

- Legacy AIX platforms (for example, single app per server, fractional CPUs)
- Web/file/print (for example, consolidation Linux, IA-32 to sub-CPU partitions)
- Application development/testing (secure, isolation)

Responsive OS images in Shared Processor LPARs with higher utilization

- Production (for example, OLTP) with batch systems
- Overlapping production systems (for example, peak of one; valley of another)
- Virtual blade environment (handle bursts of activity)

Exploitation of Virtual I/O, VLAN

- Multi-tier applications – (consolidation of Web/application/database)
- Using virtual disk to consolidate relatively small boot images
- Consolidation of low-utilized Ethernet adapters

Figure 3-50. Scenarios – Where APV is Likely to Have Significant Impact AM711.0

Notes:

Take it to the Next Level… Ask Clients:

How can you use the VALUE of IBM System p servers?

Think about…How to optimize “UNIX or Linux infrastructure”?
How can you leverage infrastructure for Business Value?

It’s not only a cost based comparison, but an investment case decision based on ROI.

Should you invest in a UNIX or Linux infrastructure that will enable:
…quicker response to business needs?
…better SLA levels?
…overall cost reduction?

How can you move to a virtualized environment to achieve these breakthrough savings?

Figure 3-51. Take it to the Next Level… Ask Clients: AM711.0

Notes:

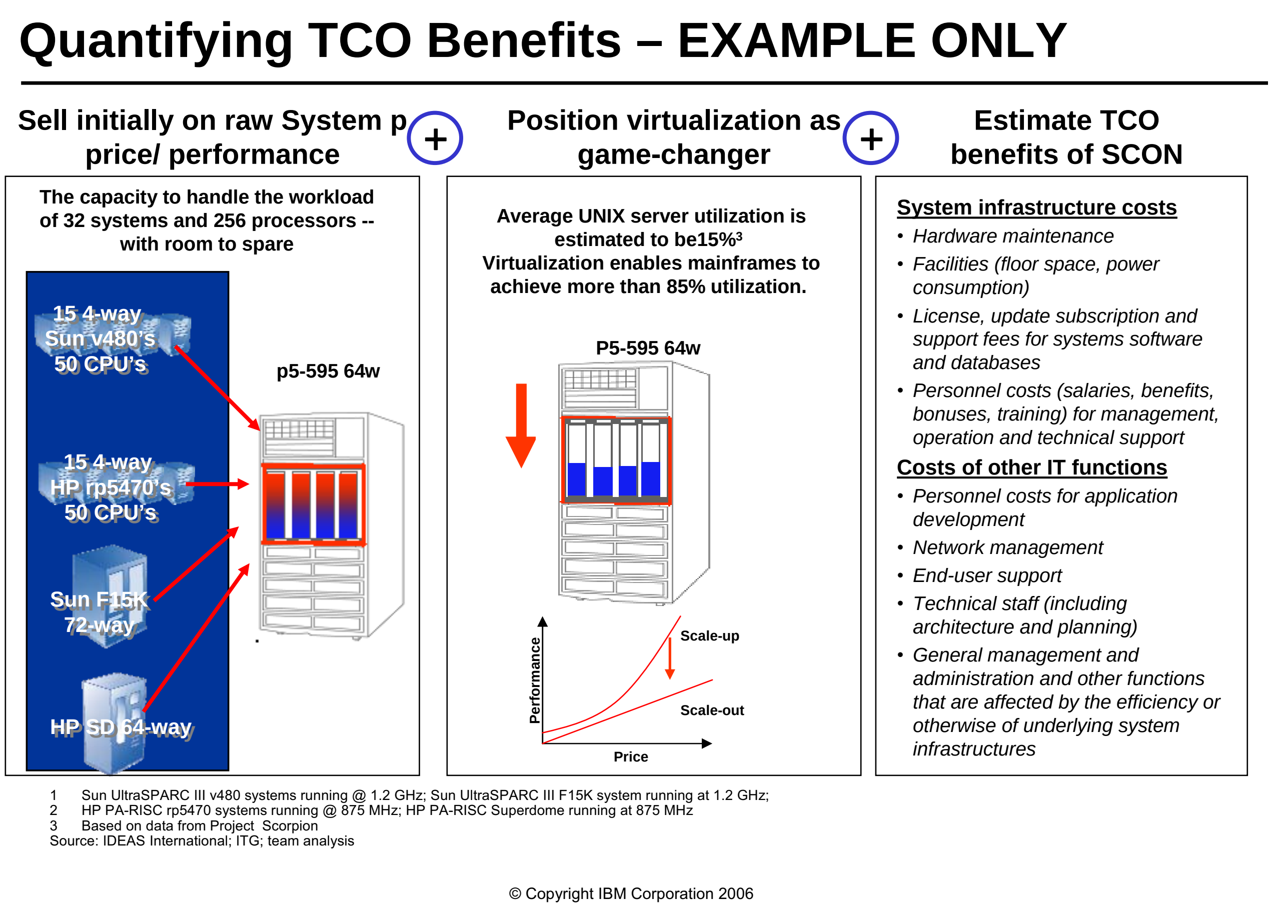

Quantifying TCO Benefits – EXAMPLE ONLY

Sell initially on raw System p price/ performance + **Position virtualization as game-changer** + **Estimate TCO benefits of SCON**

The capacity to handle the workload of 32 systems and 256 processors -- with room to spare

Average UNIX server utilization is estimated to be15%³
Virtualization enables mainframes to achieve more than 85% utilization.

System infrastructure costs

- *Hardware maintenance*
- *Facilities (floor space, power consumption)*
- *License, update subscription and support fees for systems software and databases*
- *Personnel costs (salaries, benefits, bonuses, training) for management, operation and technical support*

Costs of other IT functions

- *Personnel costs for application development*
- *Network management*
- *End-user support*
- *Technical staff (including architecture and planning)*
- *General management and administration and other functions that are affected by the efficiency or otherwise of underlying system infrastructures*

1 Sun UltraSPARC III v480 systems running @ 1.2 GHz; Sun UltraSPARC III F15K system running at 1.2 GHz;
2 HP PA-RISC rp5470 systems running @ 875 MHz; HP PA-RISC Superdome running at 875 MHz
3 Based on data from Project Scorpion
Source: IDEAS International; ITG; team analysis

Figure 3-52. Quantifying TCO Benefits – EXAMPLE ONLY AM711.0

Notes:

The p5-595 is an ideal platform for server consolidation because of the tremendous performance that the system delivers. With the performance and scalability the p5-595 can now allow our consolidate the vast majority of their infrastructure onto a single platform. In this example, a client could consolidate everything listed on this slide to JUST ONE p5-595 64-way system:

15 Sun V480s 1.2GHz 4-way systems

1 Sun F15K 1.2GHz 72-way system

15 HP HP PA-RISC rp5470 875MHZ 4-way systems

1 HP PA-RISC Superdome 875MHz 64-way systems

Putting 32 systems with 256 processors onto a single p5-595 64-way system, providing tremendous financial savings as well as savings in flexibility and manageability.

Web Application Server Example

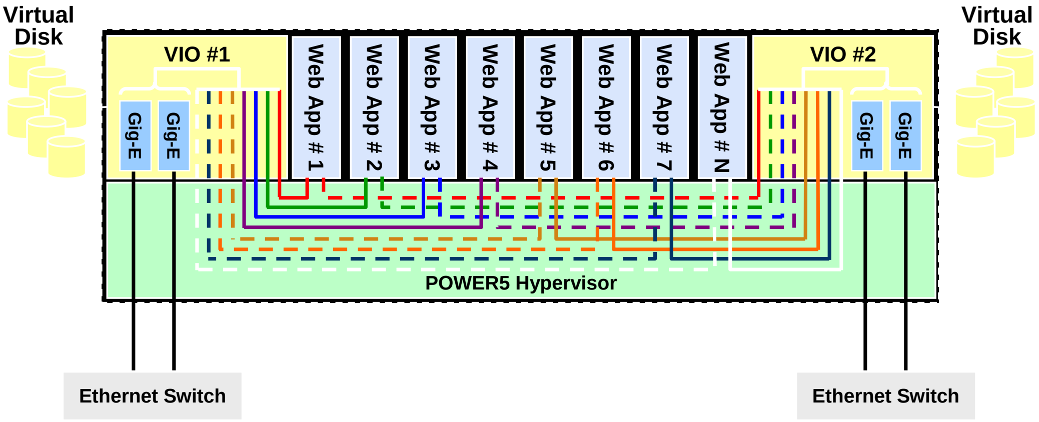

- Reduce costs
  - **Share resources between fluctuating workloads**
  - **Potentially reduce software costs**
  - **Lower network adapter and port costs**
- Improve Service Levels
  - **Resize instances to match capacity requirements**
- Improve systems management
  - **Simplify capacity management**
  - **Clone new instances easily**
- Reduce time to market
  - **Deploy new instances rapidly**

Figure 3-53. Web Application Server Example AM711.0

Notes:

High Availability Example

- Reduce cost
  - Significantly reduce idle high availability hardware
  - Automatically extend HA LPARs when needed
- Cross campus disaster recovery
  - Combine HA / Test and DR
  - Significantly reduce disaster recovery time over DR hosting site

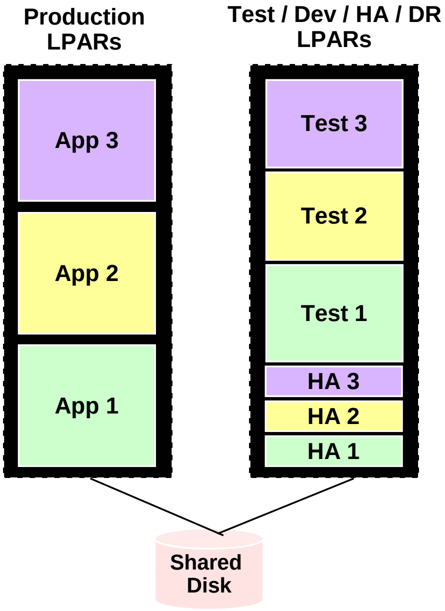

Figure 3-54. High Availability Example AM711.0

Notes:

Mixed Production and Test Environment

Data Warehouse Customer Example

- Reduced costs
  - Reduced the number of physical machines to manage
  - Used idle test and development resources for production data warehouse
- Improved service levels
  - Eliminated older less reliable hardware
  - Significantly reduced the nightly batch window without additional hardware
  - Utilized reserve capacity on demand for month or year end processing

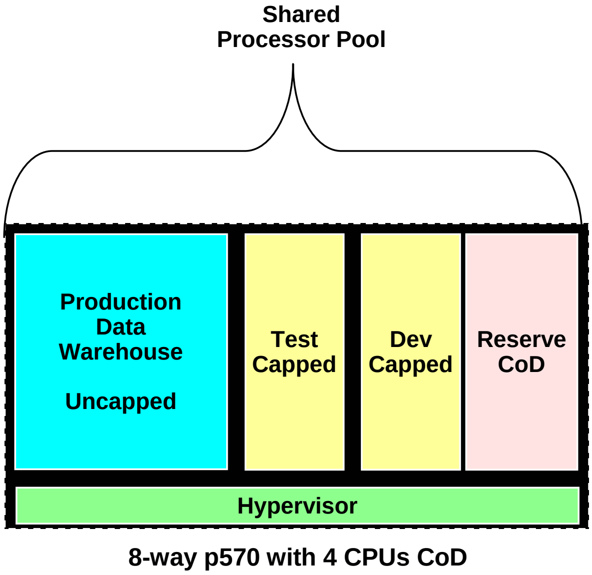

Figure 3-55. Mixed Production and Test Environment Data Warehouse Customer Example AM711.0

Notes:

Mixed Workload - Server Consolidation Example

- Reduce costs
  - Consolidate underutilized servers
  - Share resources between workloads
  - Reduce floor space costs
- Respond to market conditions
  - Rapidly provision / de-provision servers to respond to business needs
  - Change server capacity to match business requirements
- Improve systems management
  - Improve testing environments
  - Ease software migrations
  - Support different software levels to match application needs
  - Simplify capacity management
  - Simply disaster recovery

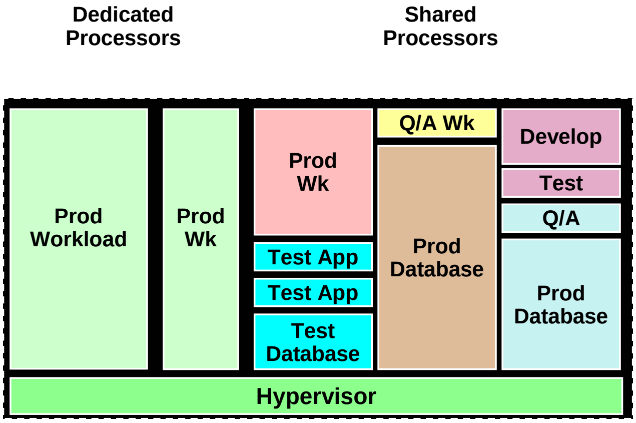

Figure 3-56. Mixed Workload - Server Consolidation Example AM711.0

Notes:

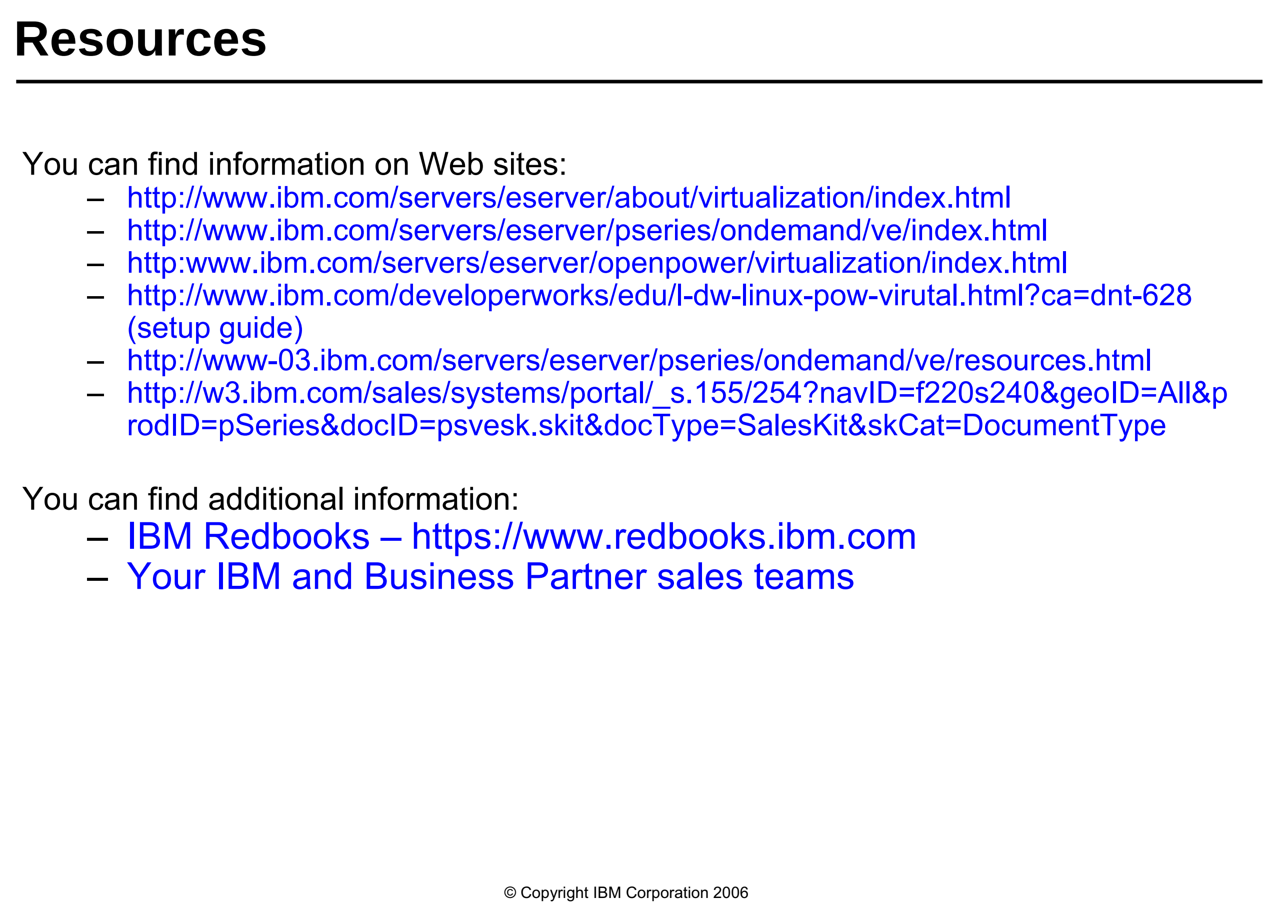

Resources

You can find information on Web sites:

- http://www.ibm.com/servers/eserver/about/virtualization/index.html
- http://www.ibm.com/servers/eserver/pseries/ondemand/ve/index.html
- http:www.ibm.com/servers/eserver/openpower/virtualization/index.html
- http://www.ibm.com/developerworks/edu/l-dw-linux-pow-virutal.html?ca=dnt-628 (setup guide)
- http://www-03.ibm.com/servers/eserver/pseries/ondemand/ve/resources.html
- http://w3.ibm.com/sales/systems/portal/_s.155/254?navID=f220s240&geoID=All&prodID=pSeries&docID=psvesk.skit&docType=SalesKit&skCat=DocumentType

You can find additional information:

- IBM Redbooks – https://www.redbooks.ibm.com
- Your IBM and Business Partner sales teams

Figure 3-57. Resources AM711.0

Notes:

Checkpoint Questions (4 of 7)

9. Partitions which have limited access to processor resources in the shared pool are_____________.
   a. Uncapped
   b. Capped
   c. Virtual
   d. None of the above

10. True or False. When more than one uncapped partition is competing for processor resources in the shared pool the preference is given to the partition with the highest priority.

11. True or False. If a partition is unable to acquire the minimum processor resources required it will not activate.

Figure 3-58. Checkpoint Questions (4 of 7) AM711.0

Notes:

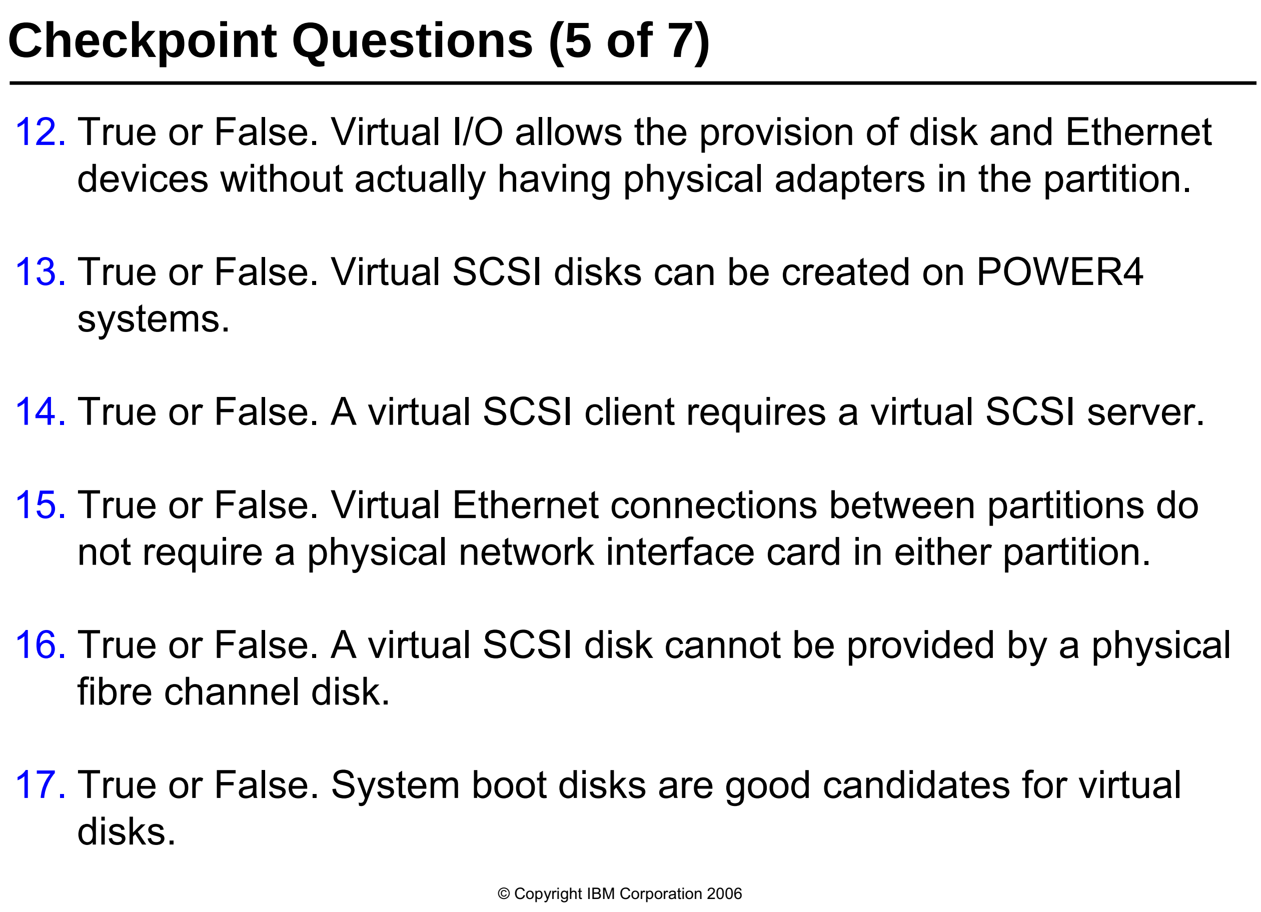

Checkpoint Questions (5 of 7)

12. True or False. Virtual I/O allows the provision of disk and Ethernet devices without actually having physical adapters in the partition.

13. True or False. Virtual SCSI disks can be created on POWER4 systems.

14. True or False. A virtual SCSI client requires a virtual SCSI server.

15. True or False. Virtual Ethernet connections between partitions do not require a physical network interface card in either partition.

16. True or False. A virtual SCSI disk cannot be provided by a physical fibre channel disk.

17. True or False. System boot disks are good candidates for virtual disks.

Figure 3-59. Checkpoint Questions (5 of 7) AM711.0

Notes:

Topic Summary

You should now be able to:

- Discuss virtualization concepts.
- Understand micro-partitioning and its role in virtualization.
- Talk about the major concepts and advantages of the IBM virtualization capabilities of server p5.

Figure 3-60. Topic Summary AM711.0

Notes:

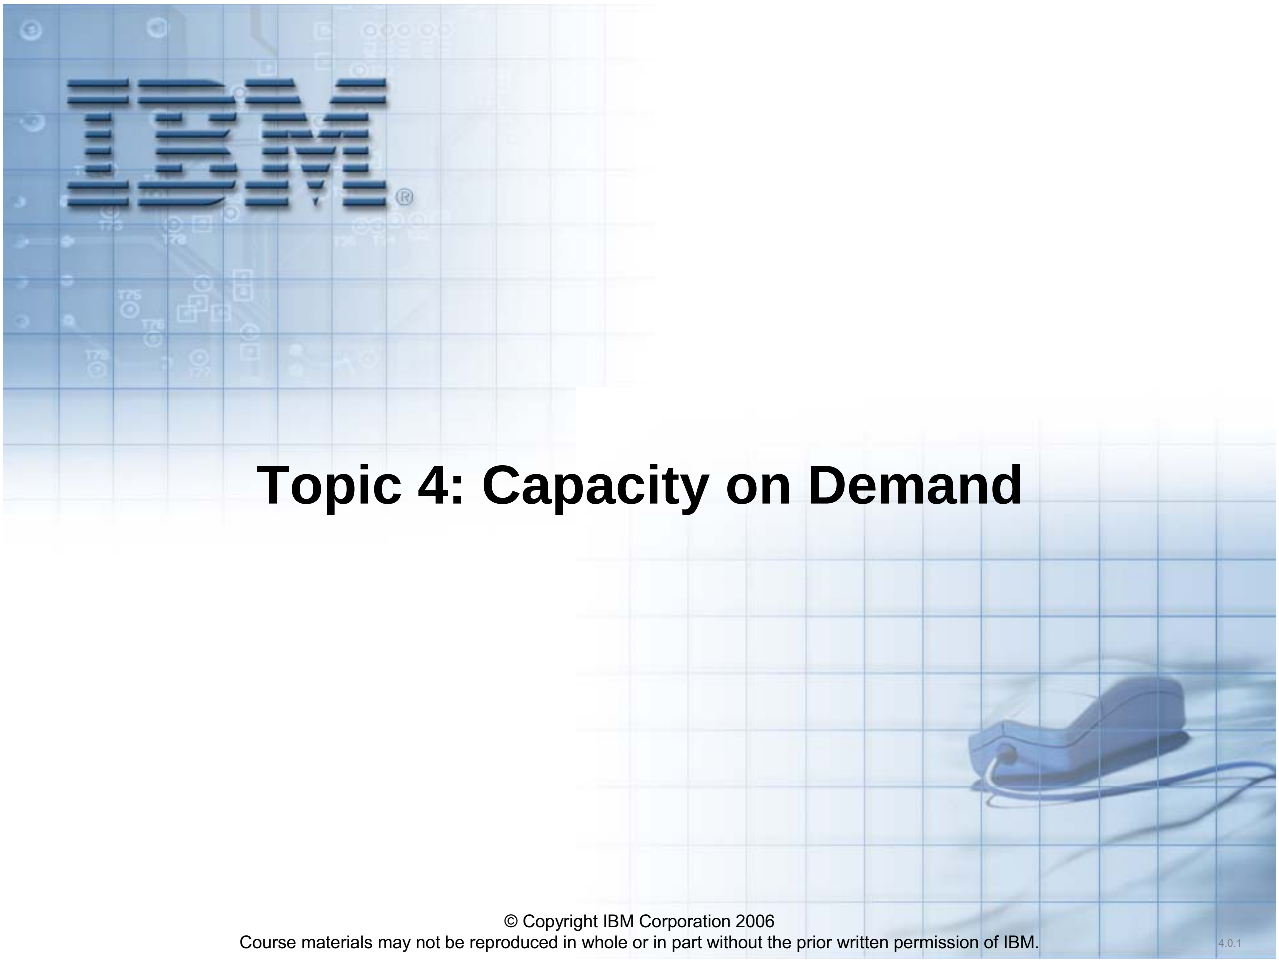

Topic 4: Capacity on Demand

Figure 3-61. Topic 4: Capacity on Demand AM711.0

Notes:

Topic Objectives

After completing this topic, you should be able to:

- Explain why Capacity on Demand features are important to your clients.
- Highlight System p Capacity on Demand abilities.
- Describe differences between pSeries and p5 Capacity on Demand features.

Figure 3-62. Topic Objectives AM711.0

Notes:

Why is Capacity on Demand Important to You?

- **Server capacity you need, when you need it to help lower** total cost of acquisition
- **Address your planned IT capacity growth needs** quickly, non-disruptively
- **Built-in flexibility to address unplanned,** temporary spikes in demand
- **Increased configuration** flexibility
- **Increased** reliability **with processor and memory sparing**
- Deploy **new** services quickly **to react to changing business requirements**
- **Gain additional workload throughput and** automated systems performance leveling

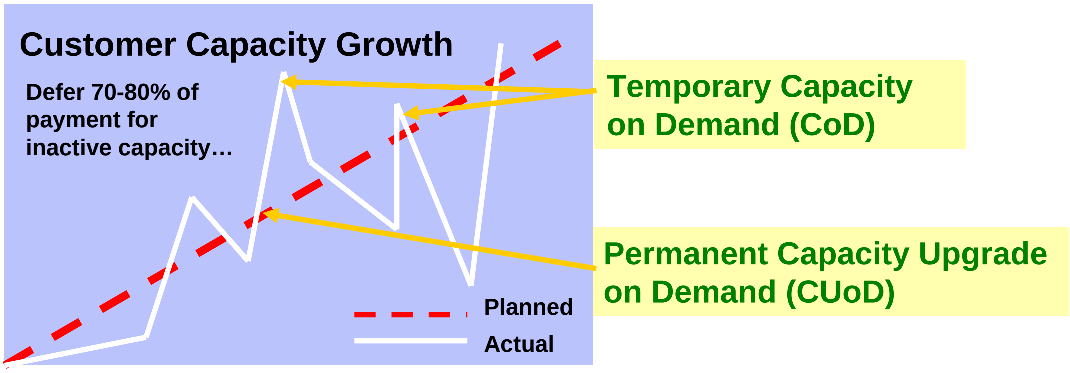

Customer Capacity Growth

Defer 70-80% of payment for inactive capacity…

Temporary Capacity on Demand (CoD)

Permanent Capacity Upgrade on Demand (CUoD)

*Capacity on Demand Features available on selected models of System p5 and eServer p5

Figure 3-63. Why is Capacity on Demand Important to You? AM711.0

Notes:

Organizations in a variety of industries have a long term growth plan and experience short-term spikes in processor utilization. With the Capacity on Demand (CoD) features optionally available on selected eServer p5 systems, clients can respond quickly and to their planned changing application workload demand by permanently activating processors or memory when planned (growth on the dashed line). This process is available immediately, and does not require a system reboot, so there is no disruption in production processing. For those temporary processing spikes in demand, there are Capacity on Demand offerings available on selected eServer p5 systems to match peak processing needs (white line with peaks and valleys). These temporary activations also have an important advantage in that they enable organizations to “activate” all the inactive CoD processors in the shared processor pool with AIX5.3, and increase total systems throughput. The bottom line is that IBM Capacity on Demand offerings allow On Demand business to react quickly to changing business conditions, simplify management and administration, whether through easier consolidation of multiple servers into a larger one, or through more automated resource availability for temporary peak capacity demands.

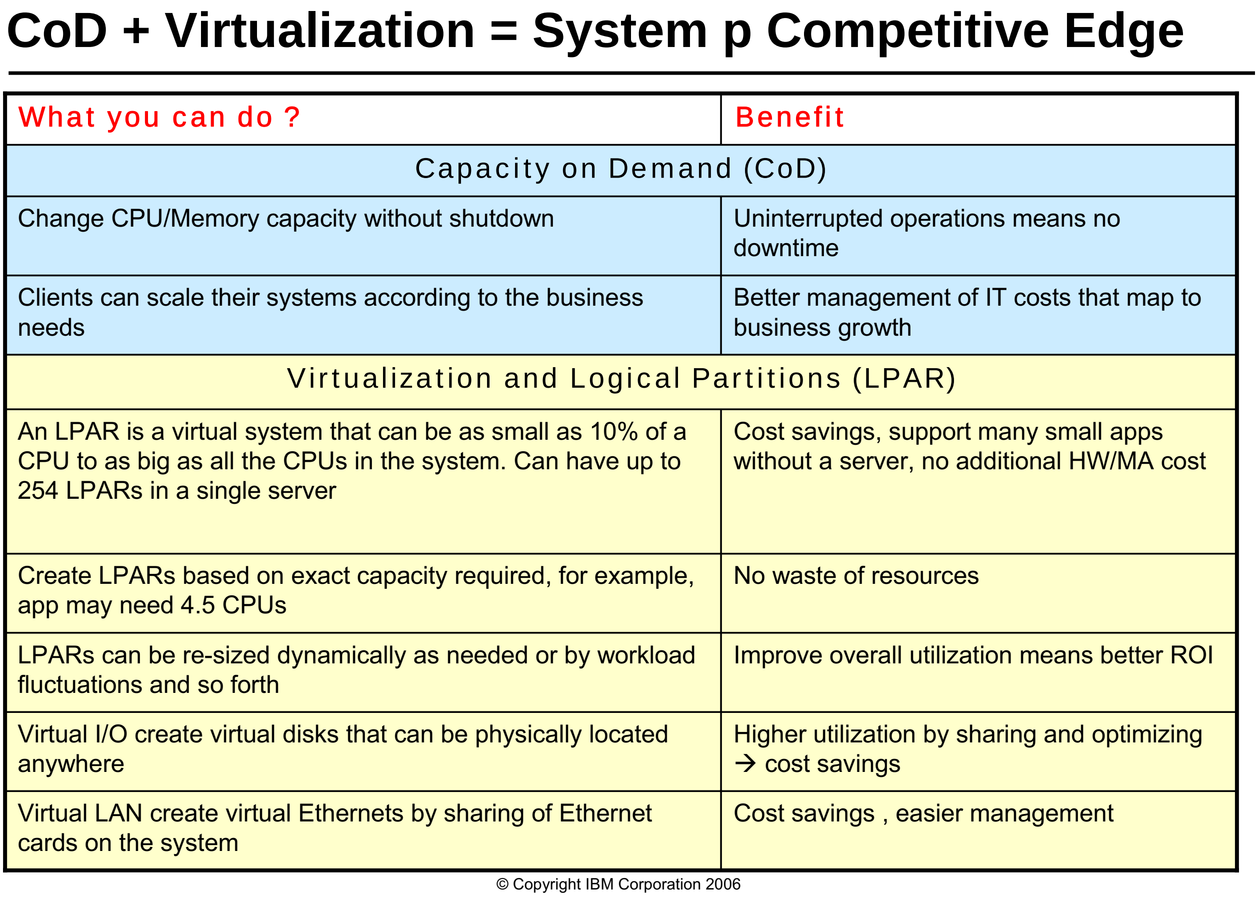

CoD + Virtualization = System p Competitive Edge

What you can do ?	Benefit
Capacity on Demand (CoD)	
Change CPU/Memory capacity without shutdown	Uninterrupted operations means no downtime
Clients can scale their systems according to the business needs	Better management of IT costs that map to business growth
Virtualization and Logical Partitions (LPAR)	
An LPAR is a virtual system that can be as small as 10% of a CPU to as big as all the CPUs in the system. Can have up to 254 LPARs in a single server	Cost savings, support many small apps without a server, no additional HW/MA cost
Create LPARs based on exact capacity required, for example, app may need 4.5 CPUs	No waste of resources
LPARs can be re-sized dynamically as needed or by workload fluctuations and so forth	Improve overall utilization means better ROI
Virtual I/O create virtual disks that can be physically located anywhere	Higher utilization by sharing and optimizing → cost savings
Virtual LAN create virtual Ethernets by sharing of Ethernet cards on the system	Cost savings , easier management

Figure 3-64. CoD + Virtualization = System p Competitive Edge AM711.0

Notes:

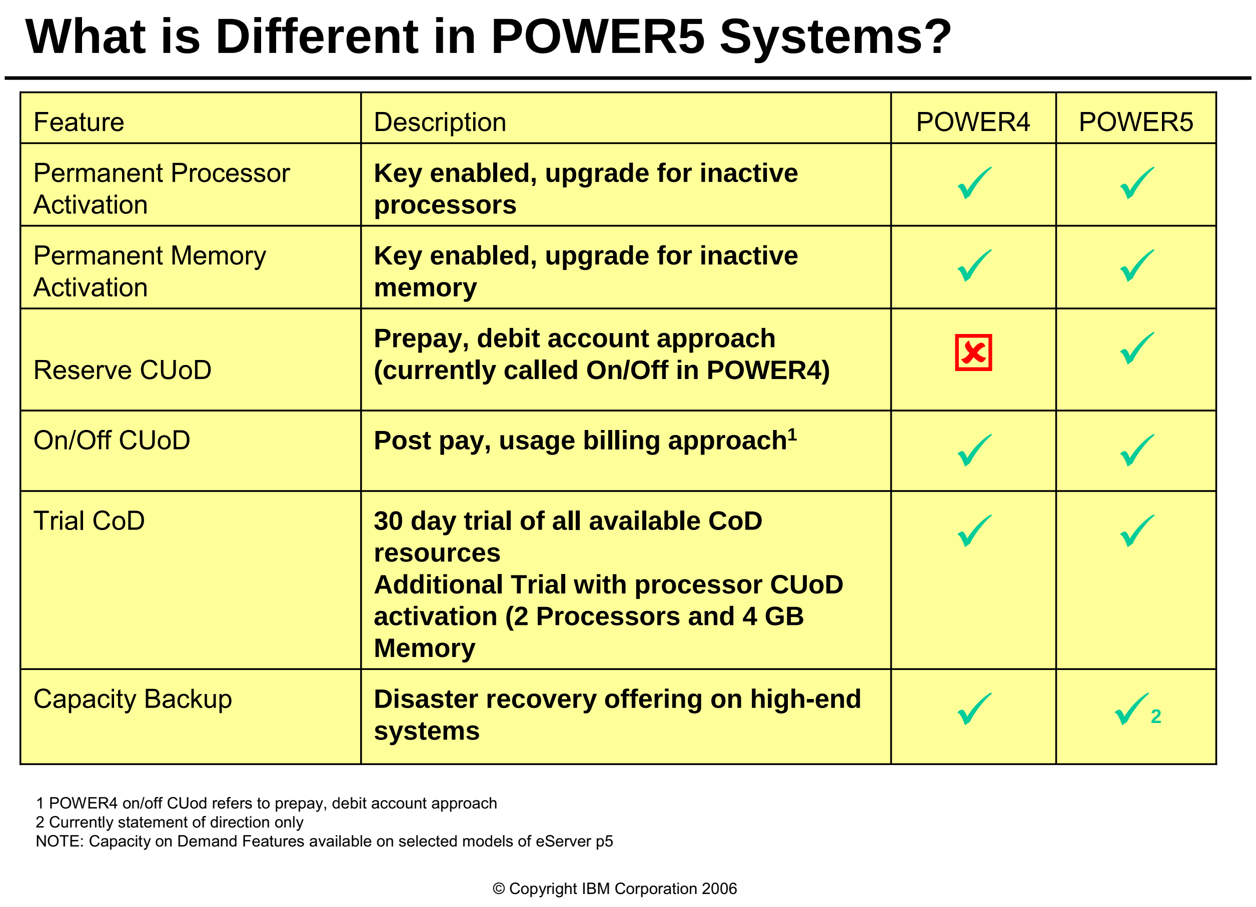

What is Different in POWER5 Systems?

Feature	Description	POWER4	POWER5
Permanent Processor Activation	**Key enabled, upgrade for inactive processors**	✓	✓
Permanent Memory Activation	**Key enabled, upgrade for inactive memory**	✓	✓
Reserve CUoD	**Prepay, debit account approach (currently called On/Off in POWER4)**	☒	✓
On/Off CUoD	**Post pay, usage billing approach**[1]	✓	✓
Trial CoD	**30 day trial of all available CoD resources** **Additional Trial with processor CUoD activation (2 Processors and 4 GB Memory**	✓	✓
Capacity Backup	**Disaster recovery offering on high-end systems**	✓	✓[2]

1 POWER4 on/off CUod refers to prepay, debit account approach
2 Currently statement of direction only
NOTE: Capacity on Demand Features available on selected models of eServer p5

Figure 3-65. What is Different in POWER5 Systems? AM711.0

Notes:

There are some changes in naming and terms and conditions between Capacity on Demand in Power4 today and Capacity on Demand in Power5 that you should be aware of. Point out that On/Off is currently a prepay 30 processor day allocation in Power4, and changes to Post pay usage billing in Power5.

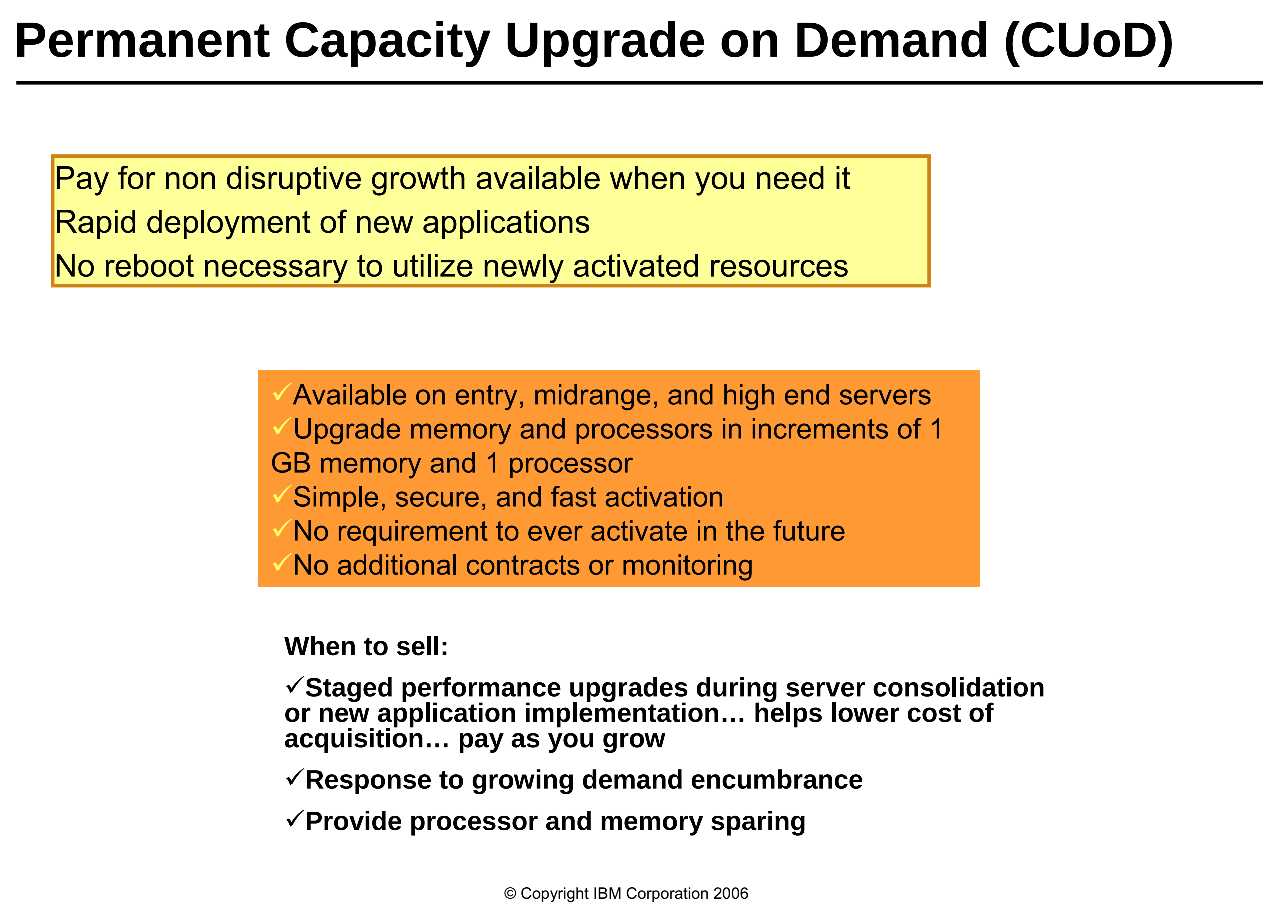

Permanent Capacity Upgrade on Demand (CUoD)

Pay for non disruptive growth available when you need it
Rapid deployment of new applications
No reboot necessary to utilize newly activated resources

✓Available on entry, midrange, and high end servers
✓Upgrade memory and processors in increments of 1 GB memory and 1 processor
✓Simple, secure, and fast activation
✓No requirement to ever activate in the future
✓No additional contracts or monitoring

When to sell:

✓Staged performance upgrades during server consolidation or new application implementation… helps lower cost of acquisition… pay as you grow

✓Response to growing demand encumbrance

✓Provide processor and memory sparing

Figure 3-66. Permanent Capacity Upgrade on Demand (CUoD) AM711.0

Notes:

Capacity Upgrade on Demand (CUoD) allows companies to purchase and activate additional permanent processor or memory capacity when the transaction load requires it. This feature helps support rapid, nondisruptive deployment of new applications and performance upgrades during server consolidations, new application bring up, and to respond to immediate business opportunities. There is no requirement to reboot the system to be able to immediately utilize the activated resources when using dynamic LPAR, eliminating costly down time. Activation is fast with clients able to obtain an activation code within 24 hours.

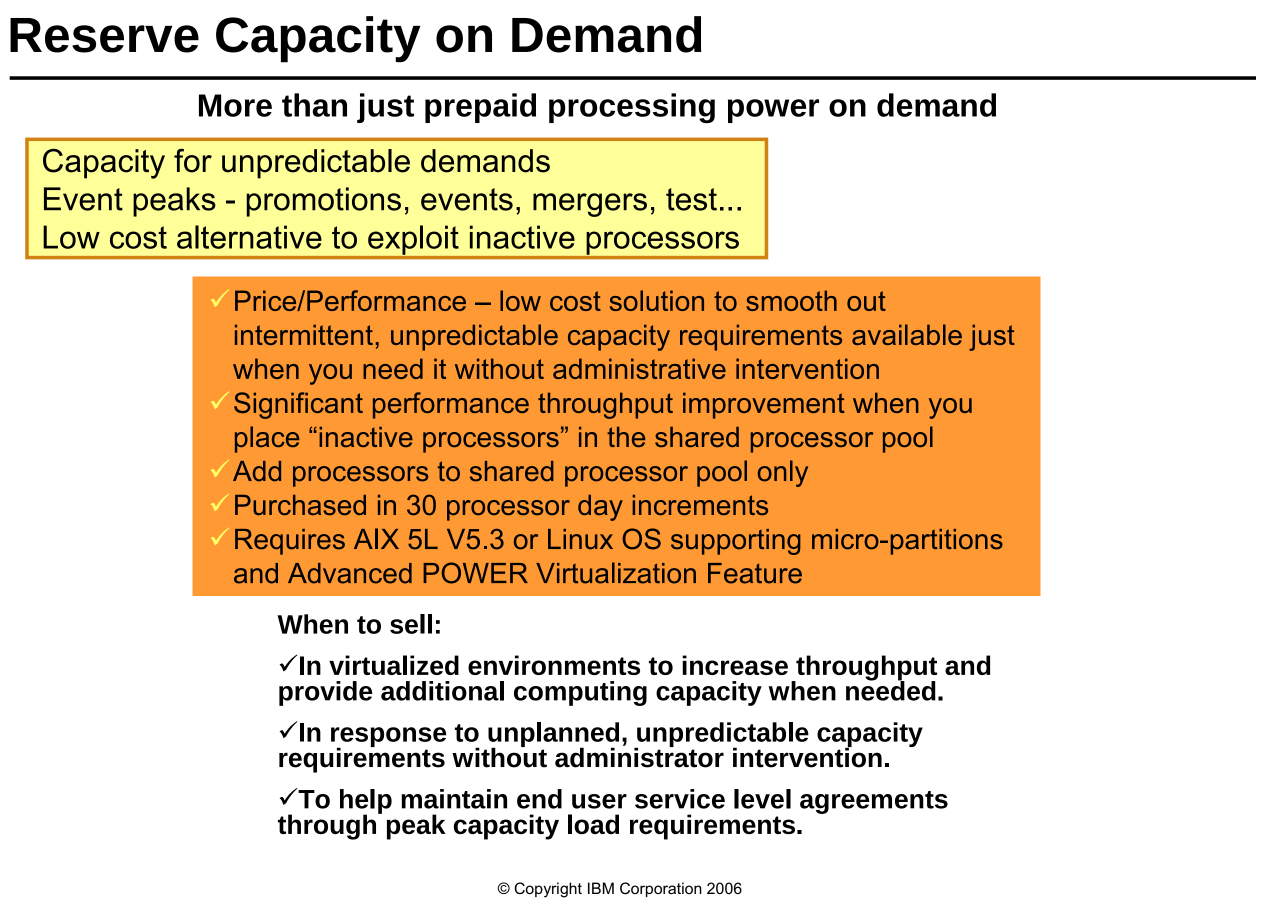

Reserve Capacity on Demand

More than just prepaid processing power on demand

Capacity for unpredictable demands
Event peaks - promotions, events, mergers, test...
Low cost alternative to exploit inactive processors

- ✓ Price/Performance – low cost solution to smooth out intermittent, unpredictable capacity requirements available just when you need it without administrative intervention
- ✓ Significant performance throughput improvement when you place “inactive processors” in the shared processor pool
- ✓ Add processors to shared processor pool only
- ✓ Purchased in 30 processor day increments
- ✓ Requires AIX 5L V5.3 or Linux OS supporting micro-partitions and Advanced POWER Virtualization Feature

When to sell:

✓In virtualized environments to increase throughput and provide additional computing capacity when needed.

✓In response to unplanned, unpredictable capacity requirements without administrator intervention.

✓To help maintain end user service level agreements through peak capacity load requirements.

Figure 3-67. Reserve Capacity on Demand AM711.0

Notes:

Reserve CoD, is designed to handle unpredictable spikes in demand, and enables organizations to activate CoD processors in response to unplanned short term workload demand or event peaks such as promotions, peaks in web server hits, or cyclical transaction processing. These processors are assigned to the shared processor pool and will automatically be used by all the shared partitions when the activation key is entered at the HMC, to provide the additional needed processing capacity without operator intervention. So long as the total system capacity remains below the level of the permanent processors in the shared processor pool, there is no deduction from the temporary processor day allocation. As peak processing requirements drive capacity requirements above the level of the paid processor capacity in the shared processor pool, the additional capacity is available automatically to handle the additional load, and processor day increments are deducted from the activation allocation as they are used. Reserve CoD provides clients a way of handling temporary capacity spikes with predictable costs and reduces management and administrative complexity for partition provisioning in a virtualized environment.

On / Off Capacity on Demand

For the predictable peeks in processing Power requirements

Repeating weekly, monthly, annual demand
Planned event peaks - promotions, events, and so forth
Temporary demand requirements

- ✓ Price/Performance – lower cost solution to cover predictable and planned peek processing requirements, available just when you need it for only as long as you need it… for memory and processors
- ✓ Activate processors and memory in dedicated or virtualized environments
- ✓ You precisely control how much resource is turned on, when it gets turned on, and when it gets released or turned off
- ✓ Data is collected by the system, transmitted to IBM and billed quarterly
- ✓ Only processor activation available on p5-550

When to sell:

✓Self managed program for temporary added memory and processors to handle additional capacity requirements.

✓For AIX 5L V5.2 or V5.3 in both dedicated and shared partition environments.

✓Help maintain end user service level agreements through peak capacity load requirements.

Figure 3-68. On / Off Capacity on Demand AM711.0

Notes:

On/Off CoD gives organizations the ability to handle planned temporary capacity requirements due to predictable business spikes. Working like a credit card with post use quarterly billing, it allows the temporary activation of reserve processors or additional memory in full-day increments as needed. When the administrator enters the enablement key, the inactive processors can be assigned to dedicated partitions, or added to the shared processor pool, the same as with Reserve CoD. If assigned to the shared processor pool, the system workload of all the shared partitions is now dispatched to all the processors in the shared processor pool. As capacity requirements increase above that of the permanent active processor level, the additional load is automatically handled, and the client is charged in full processor day increments and memory day increments until the administrator deactivates the resources at the HMC. Processor and memory day charges are incurred as soon as processors and memory are activated, and end when the resources are deactivated. The Enablement key will expire after 360 Processor Days. Resources will not be able to be activated until another enablement key is obtained from the IBM Web site.

Trial Capacity on Demand

One time, 30 days of no additional charge use
Test out inactive capacity

- ✓30 processor day trial period for all the “inactive” memory and processors installed on the server
- ✓No charge, just get the activation key from the IBM web site
- ✓You can control the amount of resource you want to test, partial upgrades, or try all that you have installed
- ✓After the test drive, then exploit one of the CoD upgrade offerings
    - ➢Permanent Activation
    - ➢Reserve CoD
    - ➢On/Off CoD

When to sell:

✓Can be used for benchmarking and evaluating use of extra resources prior to purchasing permanent or temporary activations

✓Can be used for a 30 day peak load requirement

✓After each permanent processor activation, an additional 30 day trial is available with 2 processors and 4 GB of memory

Figure 3-69. Trial Capacity on Demand AM711.0

Notes:

Trial Capacity on Demand lets clients take a test drive and see how the additional resources will improve performance and throughput. It is available immediately without the need for an activation key. The administrator can go to the HMC and enable the trial which can activate up to all the inactive CoD processors and memory available on the server. The trial is for 30 contiguous days. When a client purchases a permanent processor activation, the trial will be reset and the client may activate a trial again, but these subsequent trials will only allow the activation of two inactive processors and 4 GB of memory. The intent is to allow clients to handle additional workload with some additional resources until they get the permanent activation key for memory or processors.

Sub-Capacity Licensing Guidelines*

	# of license entitlements
Dedicated LPARs	- Based on **allocated CPUs** - If allocations dynamic, use **peak allocations**
Shared processor LPARs (capped)	- Based on **maximum entitlement** (processing units) - **Round up** if needed (no sub-CPU licensing)
Shared processor LPARs (uncapped)	- Based on **virtual processor** - **Cannot exceed** the total number of physical processors in the **shared pool** for any given application
CoD implications	- For **Permanent** activation, requires **additional** license entitlements - For **Reserve** CoD, the customer **prepays** for 30 days of both processors and OS entitlements - For **On/Off** CoD, **IBM bills** the customer for both processor days and OS license entitlement days based on the customer's reported usage - For **Trial** CoD, OS license entitlements are **included** in the trial

* All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Figure 3-70. Sub-Capacity Licensing Guidelines* AM711.0

Notes:

Checkpoint Questions (6 of 7)

18. True or False. Reserve Capacity on Demand allows activation of processor resources in a shared pool.

19. True or False. On/Off Capacity on Demand is a post pay billable approach to Capacity on Demand which is available for both POWER4 and POWER5 systems.

20. To accommodate a client's planned growth in processing demand which of the following options would be suggested.
    a. Permanent Capacity on Demand
    b. Reserve Capacity on Demand
    c. On/Off Capacity on Demand
    d. Trial Capacity on Demand

21. True or False. Trial Capacity on Demand can be activated multiple times provided the total time does not exceed the maximum entitlement.

Figure 3-71. Checkpoint Questions (6 of 7) AM711.0

Notes:

Topic Summary

You should now be able to:

- Explain why Capacity on Demand features are important to your clients.
- Highlight System p Capacity on Demand abilities.
- Describe differences between pSeries and p5 Capacity on Demand features.

Figure 3-72. Topic Summary AM711.0

Notes:

Checkpoint Questions (7 of 7)

22. True or False. Any technology improvement will boost performance of any client solution.

23. True or False. The application of technology in a creative way to solve clients’ business problems is one definition of innovation.

24. A client’s satisfaction with your solution can be enhanced by which of the following?
    a. Setting expectations appropriately.
    b. Applying technology appropriately.
    c. Communicating the benefits of the technology to the client.
    d. All of the above.

25. Which of the following are available with POWER5 architecture?
    a. Simultaneous Multi-Threading
    b. Micro-Partitioning
    c. Dynamic power management
    d. All of the above

26. True or False. Simultaneous multi-threading is the same as hyperthreading, IBM just gave it a different name.

Figure 3-73. Checkpoint Questions (7 of 7) AM711.0

Notes:

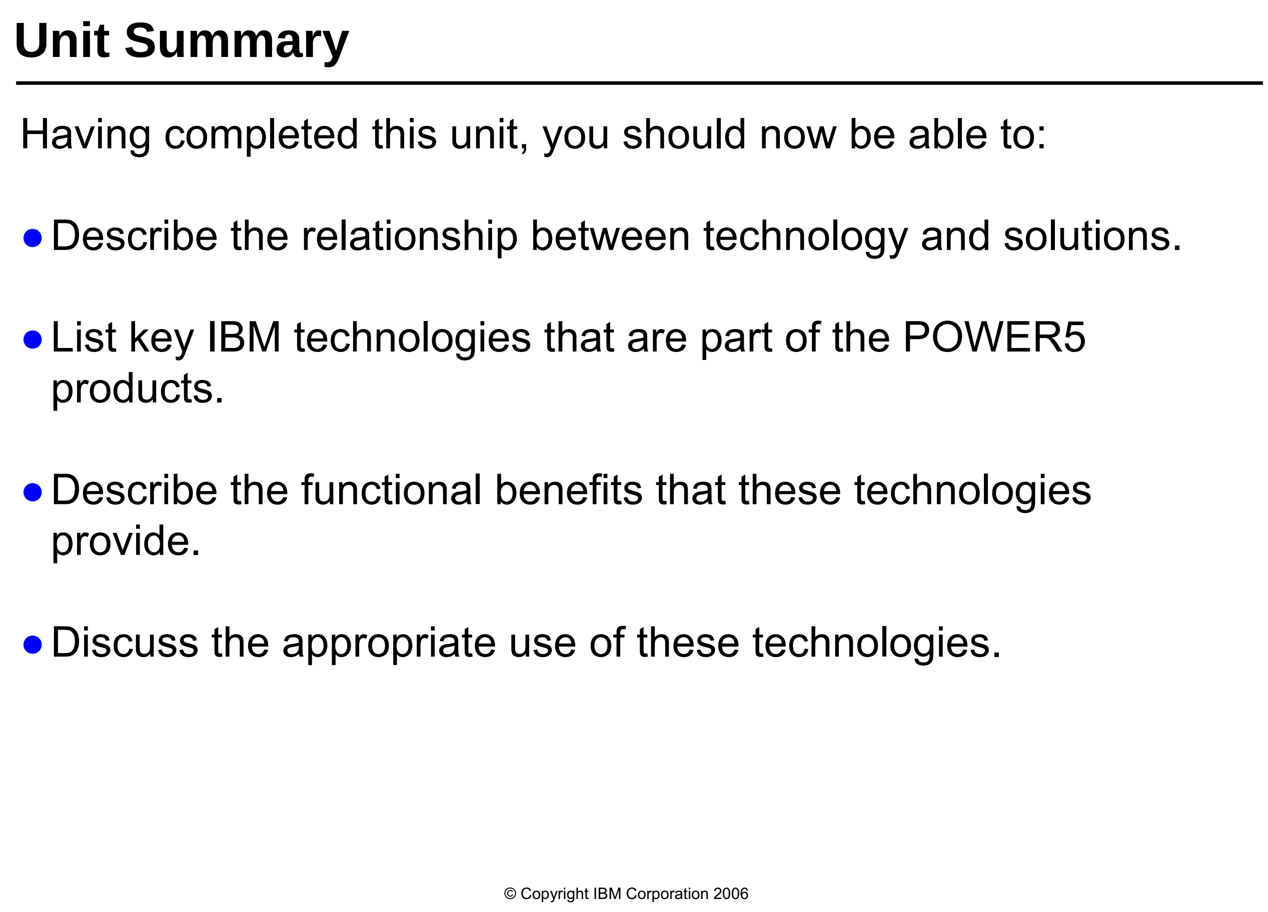

Unit Summary

Having completed this unit, you should now be able to:

- Describe the relationship between technology and solutions.
- List key IBM technologies that are part of the POWER5 products.
- Describe the functional benefits that these technologies provide.
- Discuss the appropriate use of these technologies.

Figure 3-74. Unit Summary AM711.0

Notes:

Any Questions?

Figure 3-75. Any Questions? AM711.0

Notes:

Unit 4. Servers

What This Unit Is About

This unit describes the IBM System p servers and related products.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Describe the advantages to the client of the IBM System p servers and related products
- Identify and position the current range of:
  - IBM System p5 servers
  - IBM BladeCenter servers
  - IBM IntelliStation POWER workstations
  - IBM System p I/O drawers and racks
- Identify and position the HMC and IVM methods of managing partitioned servers

How You Will Check Your Progress

Accountability:

- Checkpoint

References

IBM System p Facts and Features

http://www-03.ibm.com/systems/p/hardware/factsfeatures.htm

Unit Objectives

At the end of this topic you should be able to:

- Describe the advantages to the client of the IBM System p servers and related products.
- Identify and position the current range of:
  - IBM System p5 servers
  - IBM IBM BladeCenter servers
  - IBM IntelliStation POWER workstations
  - IBM System p I/O drawers and racks
- Identify and position the HMC and IVM methods of managing partitioned servers.

Figure 4-1. Unit Objectives AM711.0

Notes:

Introduction

- How much do I need to know about the products?
  - Depends on your role
  - Depends on your clients
- The product range is changing, how do I keep up?
  - Facts and features
  - Announcement letters
  - Redbooks
  - This course

Figure 4-2. Introduction AM711.0

Notes:

This course is a good place to start!

Products and Technology

- Products can be thought of as packaged technologies.
  - We have already spent quite a bit of time learning about the technology
  - This section builds on the previous one
  - If you have an understanding of the technology then it is easier to handle the products

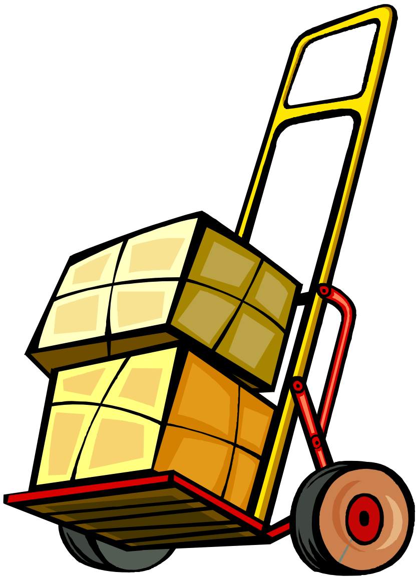

Figure 4-3. Products and Technology AM711.0

Notes:

Selling the Technology

To enable companies to achieve business breakthroughs and become an On Demand Business with IBM Systems and IBM TotalStorage, IBM has an all-encompassing commitment to openness, delivering…

Technology Innovation That Matters

Figure 4-4. Selling the Technology AM711.0

Notes:

Buzzwords, Acronyms, and TLAs

- GA
  - General Availability
- EOM
  - End of Marketing
- EOS
  - End of Support
- EOL
  - End of Life
- RFA
  - Release for Announcement

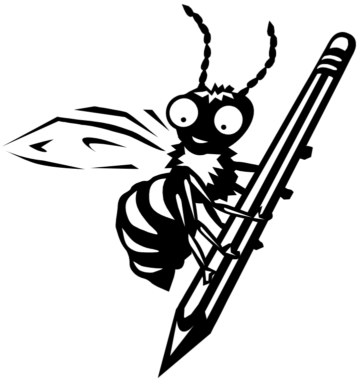

Figure 4-5. Buzzwords, Acronyms, and TLAs AM711.0

Notes:

Before we begin let us look at some of those buzzwords.

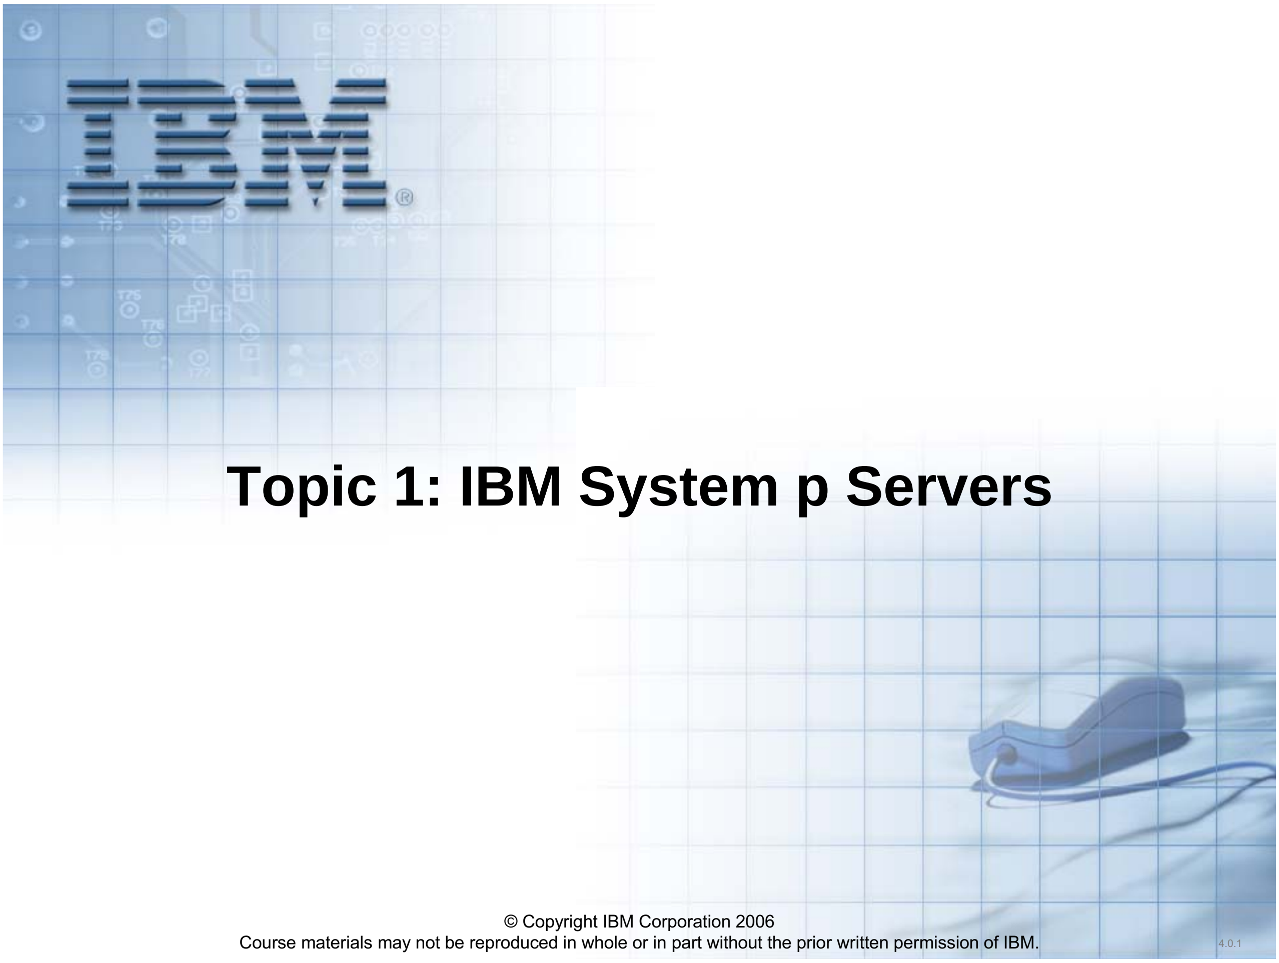

Topic 1: IBM System p Servers

4.0.1

Figure 4-6. Topic 1: IBM System p Servers AM711.0

Notes:

IBM Systems

- System Z9
  - Mission-critical mainframe data and transaction servers for large enterprises
- System i5
  - Integrated, high-performance business servers
- System p5
  - Entry, midrange, and high-end UNIX servers for commercial and technical environments
- System x
  - Intel technology based servers
- IBM BladeCenter
  - High density servers
- IBM Storage and TotalStorage
  - Disks and tapes
  - Storage area networks (SAN)
  - Storage software

Figure 4-7. IBM Systems AM711.0

Notes:

System Z9 offers mission-critical mainframe data and transaction servers for large enterprises.

System i5 servers are high-performance business servers intended for midmarket companies; highly integrated and highly available.

System x products deliver Intel technology-based servers that provide high performance in a smaller, scalable package for small- business and departmental computing.

IBM BladeCenter units gather computing resources into cost-effective, high-density spaces. The units increase the physical density of servers, improve the overall management of large numbers of servers, and share power, cooling and networking functions.

IBM Storage and TotalStorage family of products is designed to deliver quality and performance, while leading the industry in open-storage networking. The TotalStorage product family consists of: Disk storage systems, Storage area networks (SANs), Tape and optical storage, and Storage software.

Focus on IBM System p5

- For this class
- For the certification test
- Check out Top Gun classes for sales training in other areas such as:
  - TotalStorage
  - Software
  - Linux

Figure 4-8. Focus on IBM System p5 AM711.0

Notes:

Why IBM System p5?

- IBM Power Architecture technology + AIX 5L and Linux operating systems
  - Drive business innovation with powerful affordable open and adaptable UNIX and Linux offerings
- IBM Virtualization Engine + Capacity on Demand
  - Increase flexibility and resource utilization
  - Lower systems and administration costs
- Leading-edge mainframe-inspired RAS features
  - For ease of management and improved availability

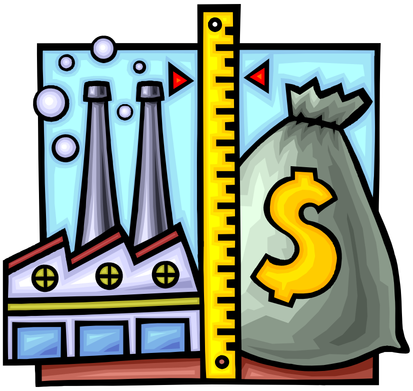

Figure 4-9. Why IBM System p5? AM711.0

Notes:

IBM Power Architecture™ technology and the AIX 5L™ and Linux® operating systems help drive business innovation with powerful, affordable, open and adaptable UNIX® and Linux offerings.

IBM Virtualization Engine™ and Capacity on Demand technologies help increase flexibility and resource utilization while helping lower systems and administration costs.

Leading-edge mainframe-inspired reliability, availability and serviceability features contribute to ease of management and improved availability.

Evolution of the IBM System p5 Express Editions

- Combining the best of pSeries and OpenPower
  - Easy to buy
  - Easy to use
  - Affordable
  - Linux AND AIX
- Offering popular, easy-to-order configurations of System p5 servers at substantial savings
  - Available in an AIX 5L Edition with reduced price AIX 5L
  - Available in an OpenPower Edition with reduced price Linux
- Allows additions to these base configurations to match up with specific application requirements

Figure 4-10. Evolution of the IBM System p5 Express Editions AM711.0

Notes:

OpenPower Edition Customers Can Get:

AIX as an option

RedHat or SUSE support

Virtualization capabilities

AIX 5L Edition Customers Can Get:

Virtualization capabilities

Linux as an option

Virtualization capabilities

IBM System p5 Express Editions

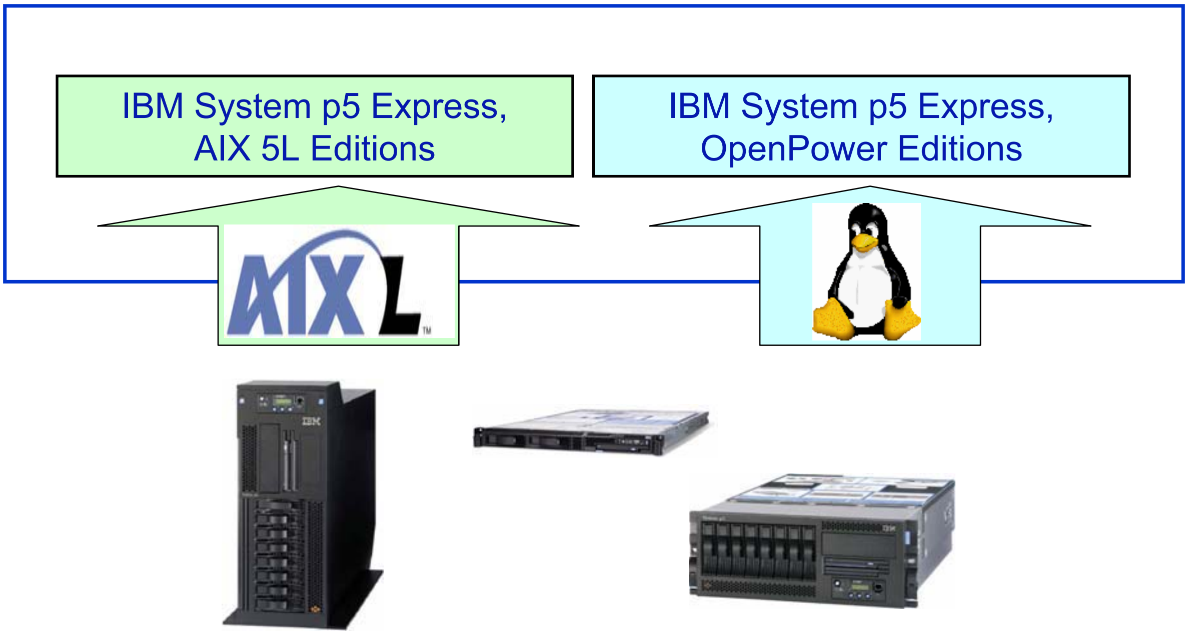

IBM System p5 Express,
AIX 5L Editions

IBM System p5 Express,
OpenPower Editions

AIX 5L Editions can add Linux and OpenPower Editions can add AIX 5L but at full pricing.

For detailed Express configurations, go to http://www-03.ibm.com/systems/p.

Figure 4-11. IBM System p5 Express Editions AM711.0

Notes:

IBM System p5 Product Line

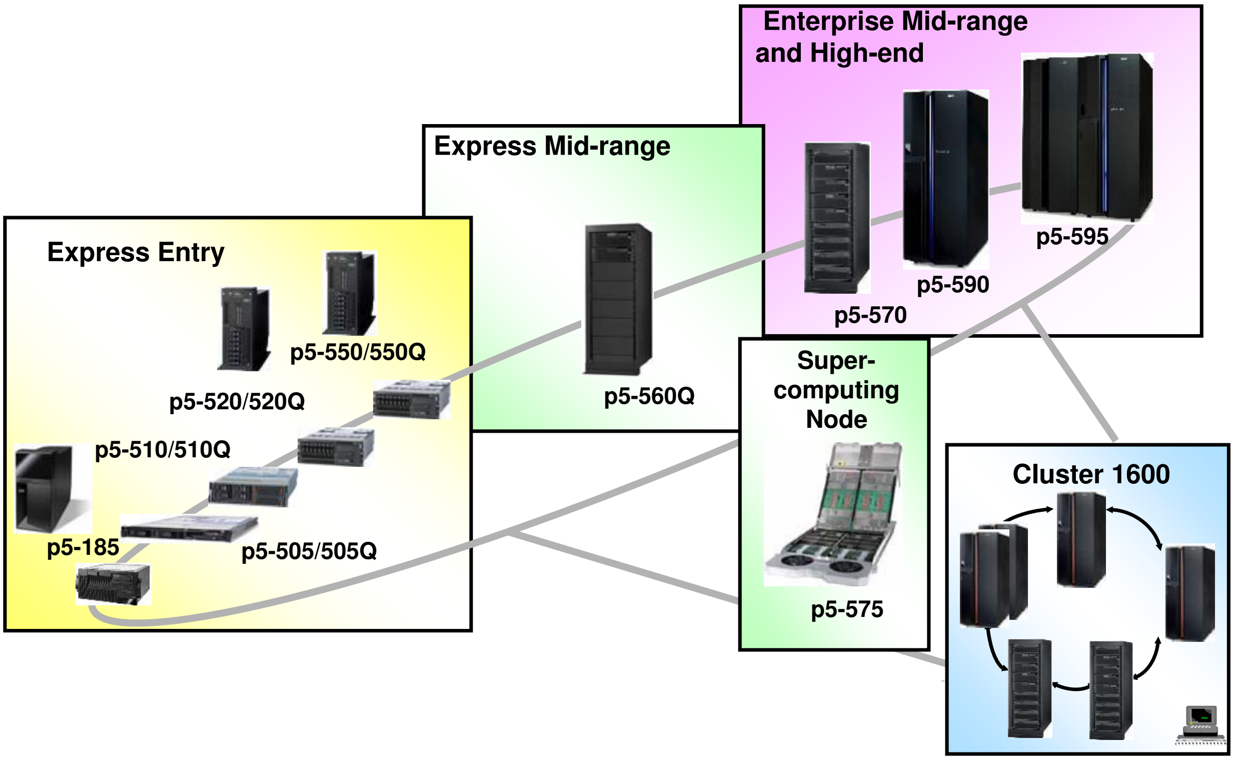

Enterprise Mid-range and High-end

p5-595

p5-590

p5-570

Express Mid-range

p5-560Q

Express Entry

p5-550/550Q

p5-520/520Q

p5-510/510Q

p5-185

p5-505/505Q

Super-computing Node

p5-575

Cluster 1600

Figure 4-12. IBM System p5 Product Line

Notes:

IBM System p5 185 Express (7037-A50)

- Exceptional value in a POWER server
- Available in deskside or rackmount
- Altivec support
- UNIX upgrade path for RS6000 150 and 170 servers
- Competitors: Dell SC 1420 Windows server

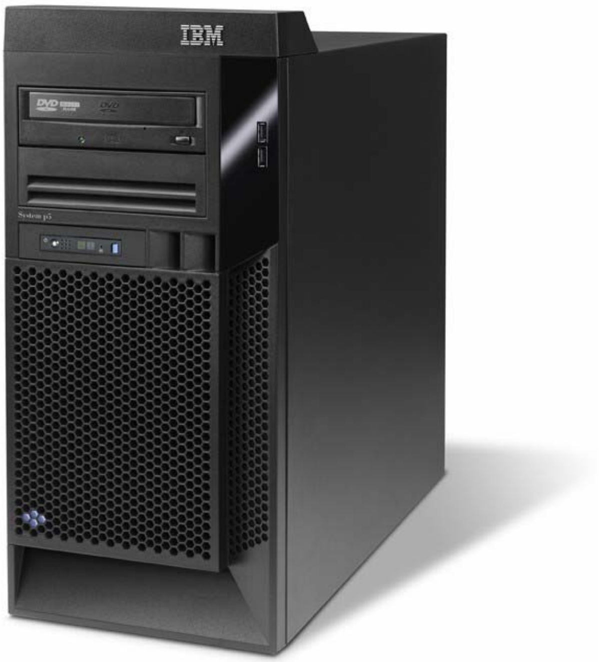

Figure 4-13. IBM System p5 185 Express (7037-A50) AM711.0

Notes:

The System p5™ 185 Express server is a small, yet powerful way to take control of your day-to-day processing needs—at a price that might surprise you.

Starting at $2,995 (IBM US list price)*

Lease for as low as $94.04/mo. for 36 months**

Processors	1- or 2-core PowerPC® 970
Clock rates (Min/Max)	2.5 GHz
System memory (Std/Max)	512MB / 8GB
Internal storage (Std/Max)	73.4GB / 1.2TB
Performance (rPerf range)***	2.48 / 4.34

IBM System p5 505 Express (9115-505)

- First POWER5 1U for space constrained applications
- Excellent HPC and infrastructure performance
- Competitive price and price performance versus Sun X4100
- Improved customer flexibility with Express offerings AIX and OpenPower

NEW!

July 2006

additional 1.9 GHz and 2.1 GHz features, 1st POWER5+ 505

Figure 4-14. IBM System p5 505 Express (9115-505) AM711.0

Notes:

As a member of the IBM System p5 Express family of servers with proven IBM POWER5™ technology, the p5-505 Express system sets a new standard for value in a high-density 1U entry UNIX® or Linux® system. It offers exceptional performance for infrastructure and HPC applications without compromising availability, flexibility or security—at the value price of an entry-level system.

Starting at $3,399 (IBM US list price)*

Lease for as low as $106.73/mo. for 36 months**

Processors	1- or 2-core POWER5/POWER5+™
Clock rates (Min/Max)	1.5 / 2.1 GHz
System memory (Std/Max)	1GB / 32GB
Internal storage (Std/Max)	73.4GB / 600GB
Performance (rPerf range)***	3.51 / 12.46

IBM System p5 505Q Express (9115-505)

- First POWER5Q 1U for space constrained applications
- Excellent HPC and infrastructure performance
- Competitive price and price performance versus Sun X4100
- Improved customer flexibility with Express offerings AIX and OpenPower

NEW!

July 2006

4-core 1.65 GHz QCM, 1U rack-mount server

Figure 4-15. IBM System p5 505Q Express (9115-505) AM711.0

Notes:

Get exceptional value from a compact server when you choose the IBM System p5 505Q Express server for database and application serving and infrastructure applications. As a member of the IBM System p5 Express family of servers with proven IBM POWER5 technology, the p5-505Q Express system sets a new standard for value in a high-density 1U entry UNIX or Linux system.

Starting at $5,505 (IBM US list price)*

Lease for as low as $172.86/mo. for 36 months**

Processors 4-core POWER5+

Clock rates (Min/Max) 1.65 GHz

System memory (Std/Max) 1GB / 32GB

Internal storage (Std/Max) 73.4GB / 600GB

Performance (rPerf range)*** 20.25

IBM System p5 510 Express (9110-51A)

The performance and scalability of POWER5+ in a 2U package

- 510: Lowest priced POWER5+ Server
- POWER5+ now in a compact 2U package
- Industry leading 1 core floating point performance[1]
- Java – Fastest 2 core server for Java business applications[1]
- Competitors Sun X4200, HP DL380, HP DL385

NEW!

July 2006

1- or 2-core 2.1 GHz, 2U rack-mount server

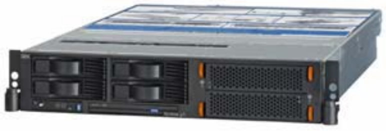

SPECjbb®2005 and SPECfp®2000 is a registered trademark of the Standard Performance Evaluation Corporation (SPEC). All hardware was tested with the SPECjbb®2005 and SPECfp®2000 benchmark suite. IBM results were compared to the top results in the www.spec.org database

Figure 4-16. IBM System p5 510 Express (9110-51A) AM711.0

Notes:

The System p5 510 Express platform is a 1- or 2-core database server with a 1.9 or 2.1 GHz IBM POWER5+ dual-core module. It can function as an entry-level application server for businesses of all sizes. The performance, reliability and affordability features of the p5-510 Express can make it a strategic platform for e-commerce application serving and compute-intensive business applications.

Starting at $4,243 (IBM US list price)*

Lease for as low as $133.23/mo. for 36 months**

Processors	1- or 2-core POWER5+
Clock rates (Min/Max)	1.9 / 2.1 GHz
System memory (Std/Max)	1GB / 32GB
Internal storage (Std/Max)	73.4GB / 1.2TB
Performance (rPerf range)***	6.11 / 12.46

IBM System p5 510Q Express (9110-51A)

- 510Q: New with 4 core QCM scalability
- Our lowest priced high performance 4 core system
- Quad core module technology enables scalability to 4 core in a 2U package
- Java – Fastest 4 core server for Java business applications[1]
- Competitors: Sun X4200, Sun T2000

NEW!

July 2006

4-core 1.65 GHz QCM, 2U rack-mount server

SPECjbb®2005 and SPECfp®2000 is a registered trademark of the Standard Performance Evaluation Corporation (SPEC). All hardware was tested with the SPECjbb®2005 and SPECfp®2000 benchmark suite. IBM results were compared to the top results in the www.spec.org database

Figure 4-17. IBM System p5 510Q Express (9110-51A) AM711.0

Notes:

The System p5 510Q Express platform is designed to meet the rigors of transaction-intensive and analytical applications. Using IBM's latest Quad-Core Module (QCM) technology with four 1.5 or 1.65 GHz POWER5+ processors, the p5-510Q Express delivers outstanding price/performance and high system throughput and is ideally suited for infrastructure application consolidation, compute-intensive business intelligence (BI) as well as scientific and engineering analytical applications.

Starting at $5,552 (IBM US list price)*

Lease for as low as $174.33/mo. for 36 months**

Processors	4-core POWER5+
Clock rates (Min/Max)	1.5 / 1.65 GHz
System memory (Std/Max)	1GB / 32GB
Internal storage (Std/Max)	73.4GB / 1.2TB
Performance (rPerf range)***	18.75

IBM System p5 520 Express (9131-52A)

- Power5+
  - 1-core 1.65 GHz SCM
  - 2-core 1.65 GHz DCM
  - 2-core 1.9 GHz DCM
- Price/performance
- Competes with Sun V40z, Sun V440

NEW!

July 2006

1- or 2-core 2.1 GHz, deskside and 4U rack-mount

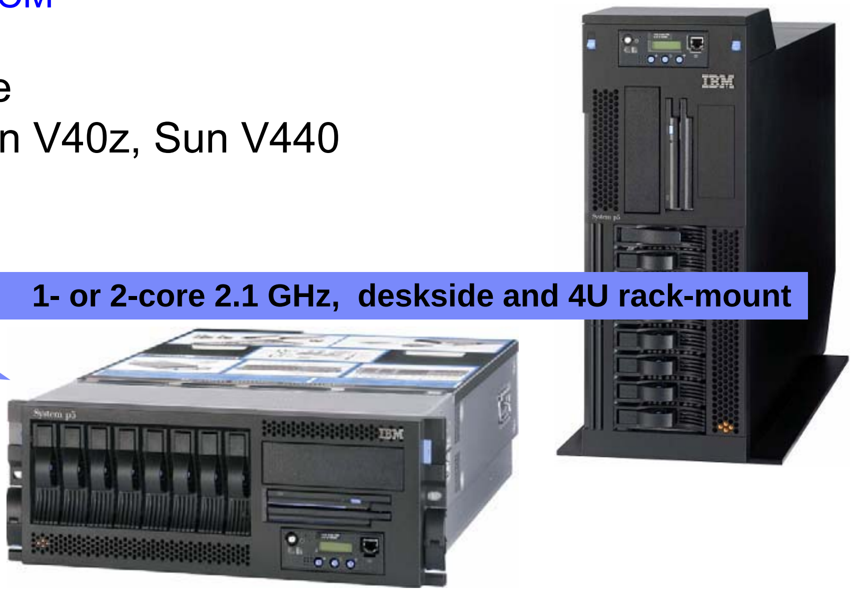

© Copyright IBM Corporation 2006

Figure 4-18. IBM System p5 520 Express (9131-52A) AM711.0

Notes:

The leadership performance and manageability position the System p5 520 Express platform with its 1.65, 1.9 or 2.1 GHz IBM POWER5+ processors for use as a small database server, a branch application server, and for highly secure e-business and business intelligence (BI) / high performance computing (HPC) applications for enterprises of all sizes.

Starting at $5,427 (IBM US list price)*

Lease for as low as $170.41/mo. for 36 months**

Processors 1- or 2-core POWER5+

Clock rates (Min/Max) 1.65 / 2.1 GHz

System memory (Std/Max) 1GB / 32GB

Internal storage (Std/Max) 73.4GB / 16.8TB (with optional I/O drawers)

Performance (rPerf range) *** 3.62 / 12.46

IBM System p5 520Q Express (9131-52A)

- **Lowest priced 4U, 4 core server available today**
- Maximum configurability in rack and desk side form factor
- Java – Fastest 4 core server for Java business applications[1]
- DLPAR standard with optional Integrated Virtualization Manager
- Competitors: Sun V40z, Sun V440

NEW!

July 2006

4-core 1.65 GHz QCM, deskside and 4U rack-mount

Our 4U, 4 core system tuned for Java and Web applications

1) Note: IBM SPECjbb data submitted to SPEC on February 14, 2006. SPECjbb®2005 is a registered trademark of the Standard Performance Evaluation Corporation (SPEC). All hardware was tested with the SPECjbb®2005 benchmark suite. IBM results were compared to the top results in the www.spec.org database as of 2/14/06.

Figure 4-19. IBM System p5 520Q Express (9131-52A) AM711.0

Notes:

The System p5 520Q Express platform is a new breed of 4-core Java™ and Web application server designed to deliver superior system throughput and configurability for demanding e-commerce and Web infrastructure applications. Powered by IBM's latest Quad-Core Module (QCM) 1.5 or 1.65 GHz POWER5+ processors, the p5-520Q Express delivers outstanding performance at compelling low prices for high transaction environments based on the Java platform or for compute-intensive database applications.

Starting at $12,610 (IBM US list price)*

Lease for as low as $395.95/mo. for 36 months**

Processors	4-core POWER5+
Clock rates (Min/Max)	1.5 / 1.65 GHz
System memory (Std/Max)	1GB / 32GB
Internal storage (Std/Max)	73.4GB / 16.8TB (with optional I/O drawers)
Performance (rPerf range)***	18.75 / 20.25

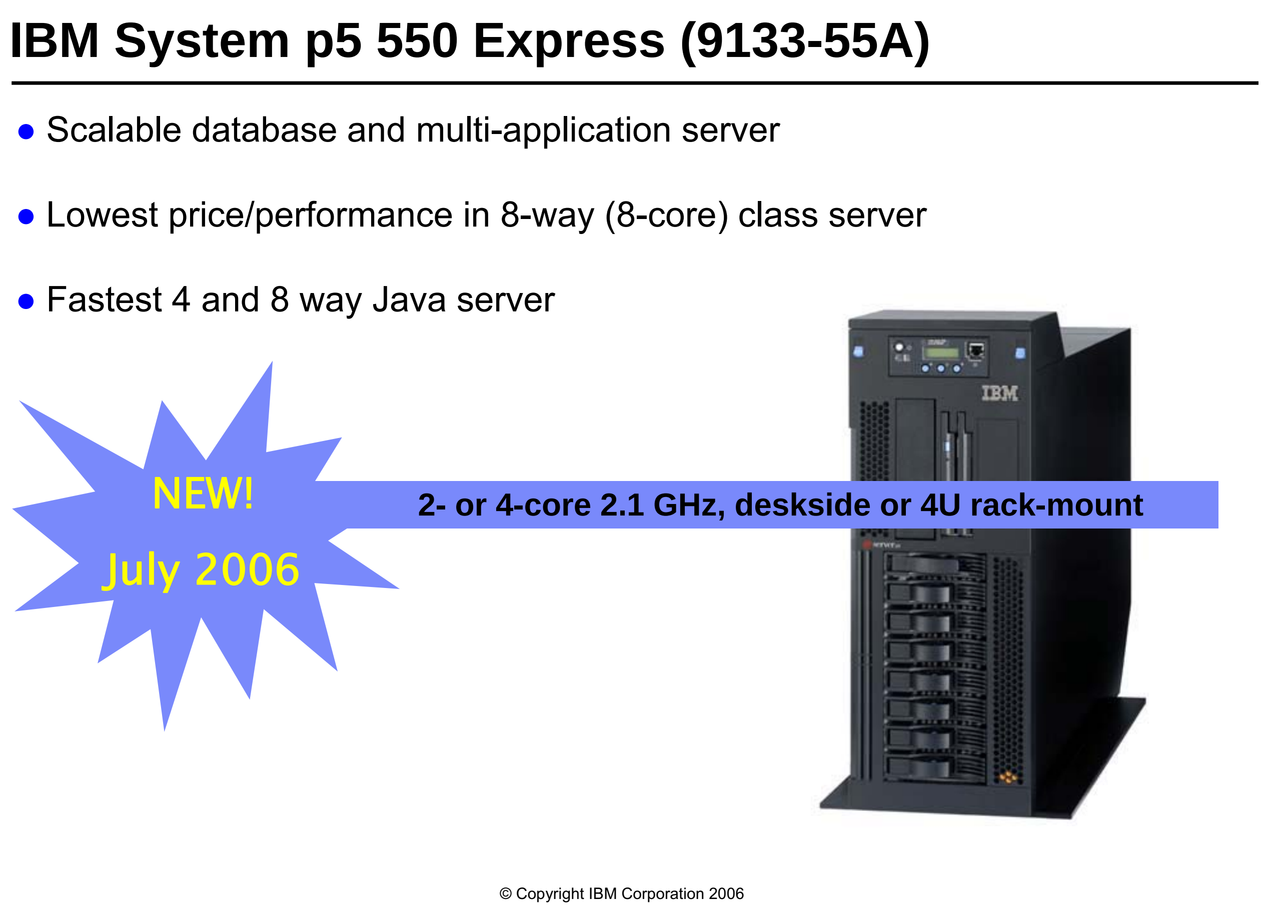

IBM System p5 550 Express (9133-55A)

- Scalable database and multi-application server
- Lowest price/performance in 8-way (8-core) class server
- Fastest 4 and 8 way Java server

NEW!

July 2006

2- or 4-core 2.1 GHz, deskside or 4U rack-mount

Figure 4-20. IBM System p5 550 Express (9133-55A) AM711.0

Notes:

The System p5 550 Express platform is a 2- or 4-core database server with 1.65, 1.9 or 2.1 GHz IBM POWER5+ dual-core modules. It can function as a scalable database or versatile departmental or regional server for businesses of all sizes. The performance, reliability and affordability features of the p5-550 Express can make it a strategic platform for scalable database servers, e-commerce application servers, Web servers, operations systems, and for BI and high performance computing (HPC) workloads.

Starting at $13,118 (IBM US list price)*

Lease for as low as $411.91/mo. for 36 months**

Processors	2- or 4-core POWER5+
Clock rates (Min/Max)	1.65 / 2.1 GHz
System memory (Std/Max)	1GB / 64GB
Internal storage (Std/Max)	73.4GB / 31.2TB (with optional I/O drawers)
Performance (rPerf range) ***	10.15 / 24.86

IBM System p5 550Q Express (9133-55A)

- Scalable database and multi-application server
- Lowest price/performance in 8-way (8-core) class server
- Fastest 4 and 8 way Java server

NEW!

July 2006

4- or 8-core 1.65 GHz QCM, deskside and 4U rack

Figure 4-21. IBM System p5 550Q Express (9133-55A) AM711.0

Notes:

The System p5 550Q Express platform is a new breed of 8-core scalable database and application server designed to meet the rigors of transaction-intensive and analytical applications. Powered by IBM's latest Quad-Core Module (QCM) 1.5 or 1.65 GHz POWER5+ processors, it is ideally suited for compute-intensive business intelligence (BI), scientific and engineering analytical applications, database applications or server consolidation.

Starting at $13,744 (IBM US list price)*

Lease for as low as $431.56/mo. for 36 months**

Processors 4- or 8-core POWER5+

Clock rates (Min/Max) 1.5 / 1.65 GHz

System memory (Std/Max) 1GB / 64GB

Internal storage (Std/Max) 73.4GB / 31.2TB (with optional I/O drawers)

Performance (rPerf range) *** 18.20 / 38.34

IBM System p5 560Q Express (9116-561)

- IBM modular ‘building block’ technology
- Leadership 16-core price/performance
- With up to 32 threads, designed to support multi-threaded HPC applications
- Dynamic logical partitioning and IBM Director standard
- Optional Advanced POWER Virtualization with Web-based Integrated Virtualization Manager

Figure 4-22. IBM System p5 560Q Express (9116-561) AM711.0

Notes:

Take back control with the easy-to-scale 4- to 16-core System p5™ 560Q server. Its unique modular building block architecture lets you add more powerful state-of-the-art IBM POWER5+™ processing capability whenever needed.

Starting at $27,109.00 (IBM US list price)*

Lease for as low as $851.22/mo. for 36 months**

Processors	4-, 8- or 16-core POWER5+
Clock rates (Min/Max)	1.5 GHz
System memory (Std/Max)	2GB / 128GB
Internal disk storage (Std/Max)	73.4GB / 3.6TB
Performance (rPerf range)***	18.75 - 65.24

IBM System p5 570 (9117-570)

First UNIX 16-core system to deliver more than 1,000,000 TPMs

- A scalable database server for applications like CRM and ERP
- Consolidation of multiple workloads and other application servers

  Best performance in its class, supporting more users per core*
  - Over 3X the performance of the HP Integrity rx8620*
- Designed with best of breed RAS features among mid-range systems
- **Choose from thousands of AIX 5L or Linux applications and IBM integrated offerings like** DB2, WebSphere, Sybase IQ & ASE, and SAS

*Based on TPC-C result submitted to tpc.org as of February 14, 2006

Figure 4-23. IBM System p5 570 (9117-570) AM711.0

Notes:

Easily scale from 2- to 16-cores with the IBM System p5™ 570. Unique IBM modular SMP architecture lets you add more powerful IBM POWER5+™ processing capability exactly when needed.

Processors 2-, 4-, 8-, 12- or 16-core POWER5+

Clock rates (Min/Max) 1.9 GHz / 2.2 GHz

System memory (Std/Max) 2GB / 512GB

Internal disk storage (Std/Max) 73.4GB / 79.2TB (with optional I/O drawers)

Performance (rPerf range)* 12.27 / 95.96

IBM System p5 575 (9118-575)

- More than 50% per node Linpack performance uplift over 8-way 575
- Scales from 2 node departmental system to 128 node supercomputer with over 2000 processors
- Designed to achieve extraordinary results in the most challenging high-performance (HPC) and business intelligence (BI) environments

Figure 4-24. IBM System p5 575 (9118-575) AM711.0

Notes:

IBM System p5 575 cluster nodes are 8- and 16-core high-powered building blocks for data mining and numeric intensive computing applications—featuring a high memory bandwidth and an ultra-dense design.

Processors	8- and 16-core IBM POWER5+
Clock rates (Min/Max)	2.2 GHz (8-core), 1.9 GHz (16-core)
System memory (Std/Max)	1GB / 256GB
Internal storage (Std/Max)	73.4GB / 2.9TB
LINPACK HPC performance (GFLOPS)1	66.44 (8-core) and 111.4 (16-core)

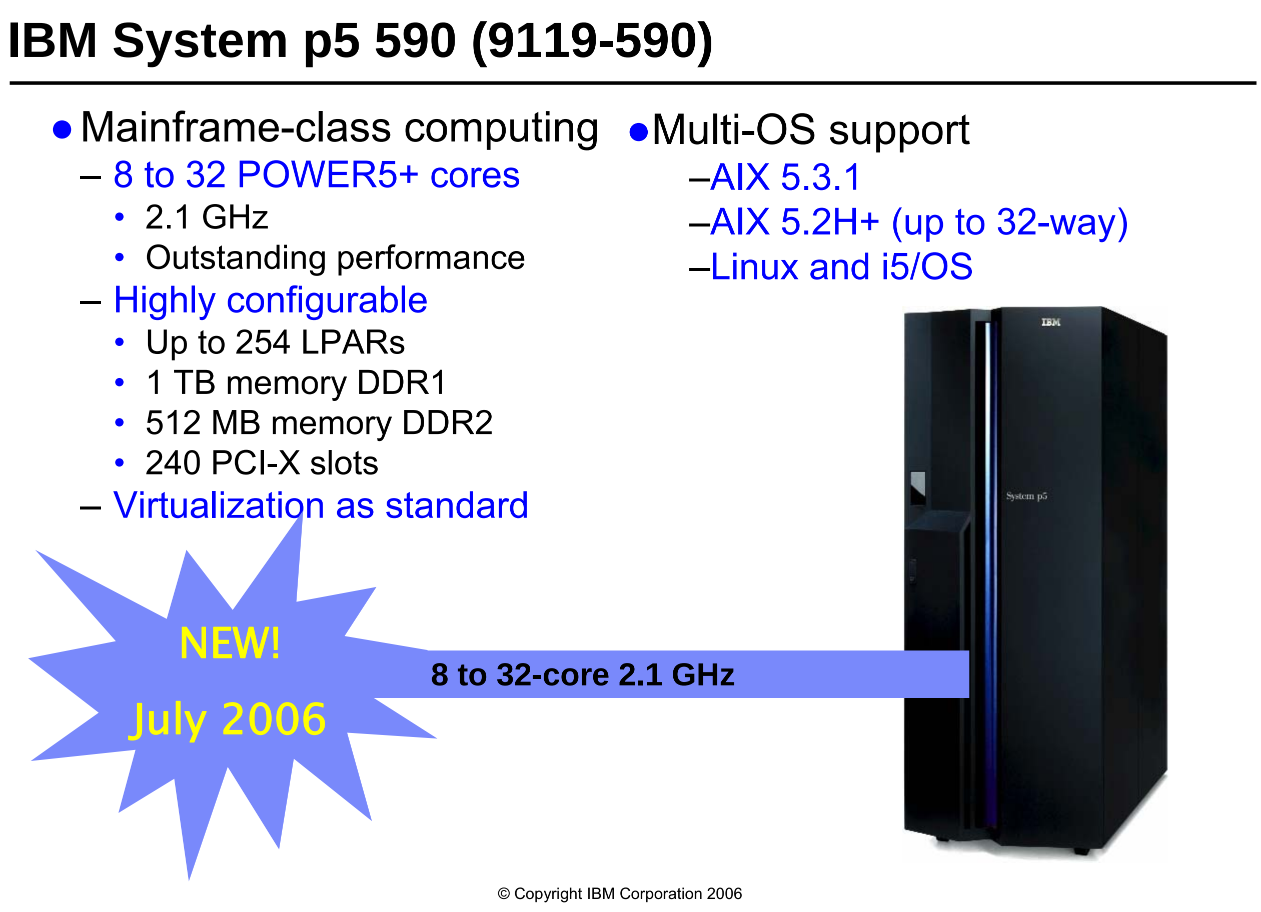

IBM System p5 590 (9119-590)

- Mainframe-class computing
  - 8 to 32 POWER5+ cores
    - 2.1 GHz
    - Outstanding performance
  - Highly configurable
    - Up to 254 LPARs
    - 1 TB memory DDR1
    - 512 MB memory DDR2
    - 240 PCI-X slots
  - Virtualization as standard
- Multi-OS support
  - AIX 5.3.1
  - AIX 5.2H+ (up to 32-way)
  - Linux and i5/OS

NEW!

July 2006

8 to 32-core 2.1 GHz

IBM
System p5

© Copyright IBM Corporation 2006

Figure 4-25. IBM System p5 590 (9119-590) AM711.0

Notes:

The IBM System p5 590 server uses fifth-generation 64-bit IBM POWER5™ technology in up to 32-core symmetric multiprocessing (SMP) configurations with Advanced POWER™ Virtualization to provide the performance, flexibility, scalability, manageability and security features needed for the consolidation of mission-critical AIX 5L™ and Linux® applications on a single system to save on hardware, software, energy and space costs.

Processors 8- to 32-core POWER5 and POWER5+

Clock rates (Min/Max) 1.65 / 2.1 GHz

System memory (Std/Max) 8GB / 1TB

Internal storage (Std/Max) 72.8GB / 18.7TB (using optional I/O drawers)

Performance (rPerf range)* 41.68 to 202.88

IBM System p5 595 (9119-595)

- Mainframe-class computing
  - Up to 64 POWER5 processors
    - Choice of 1.65, or 1.9 GHz
  - Outstanding performance
  - Highly configurable
    - Up to 254 LPARs
    - 2 TB memory DDR1
    - 1 TB memory DDR2
    - 240 PCI-X slots
  - Virtualization as standard
- Multi-OS support
  - AIX 5.3.1
  - AIX 5.2H+ (up to 32-way)
  - Linux and i5/OS

NEW!

July 2006

16 to 64-core 2.1 GHz and 2.3 GHz

Figure 4-26. IBM System p5 595 (9119-595) AM711.0

Notes:

The IBM System p5 595 server uses fifth-generation 64-bit IBM POWER5 technology in up to 64-core symmetric multiprocessing (SMP) configurations with Advanced POWER Virtualization and offers the performance, flexibility, scalability, manageability and security features needed for the consolidation of mission-critical AIX 5L and Linux applications on a single system to save on hardware, software, energy and space costs.

Processors	16- to 64-core POWER5 and POWER5+
Clock rates (Min/Max)	1.65 / 2.3 GHz
System memory (Std/Max)	8GB / 2TB
Internal storage (Std/Max)	72.8GB / 28.1TB (using optional I/O drawers)
Performance (rPerf range)*	80.96 to 393.55

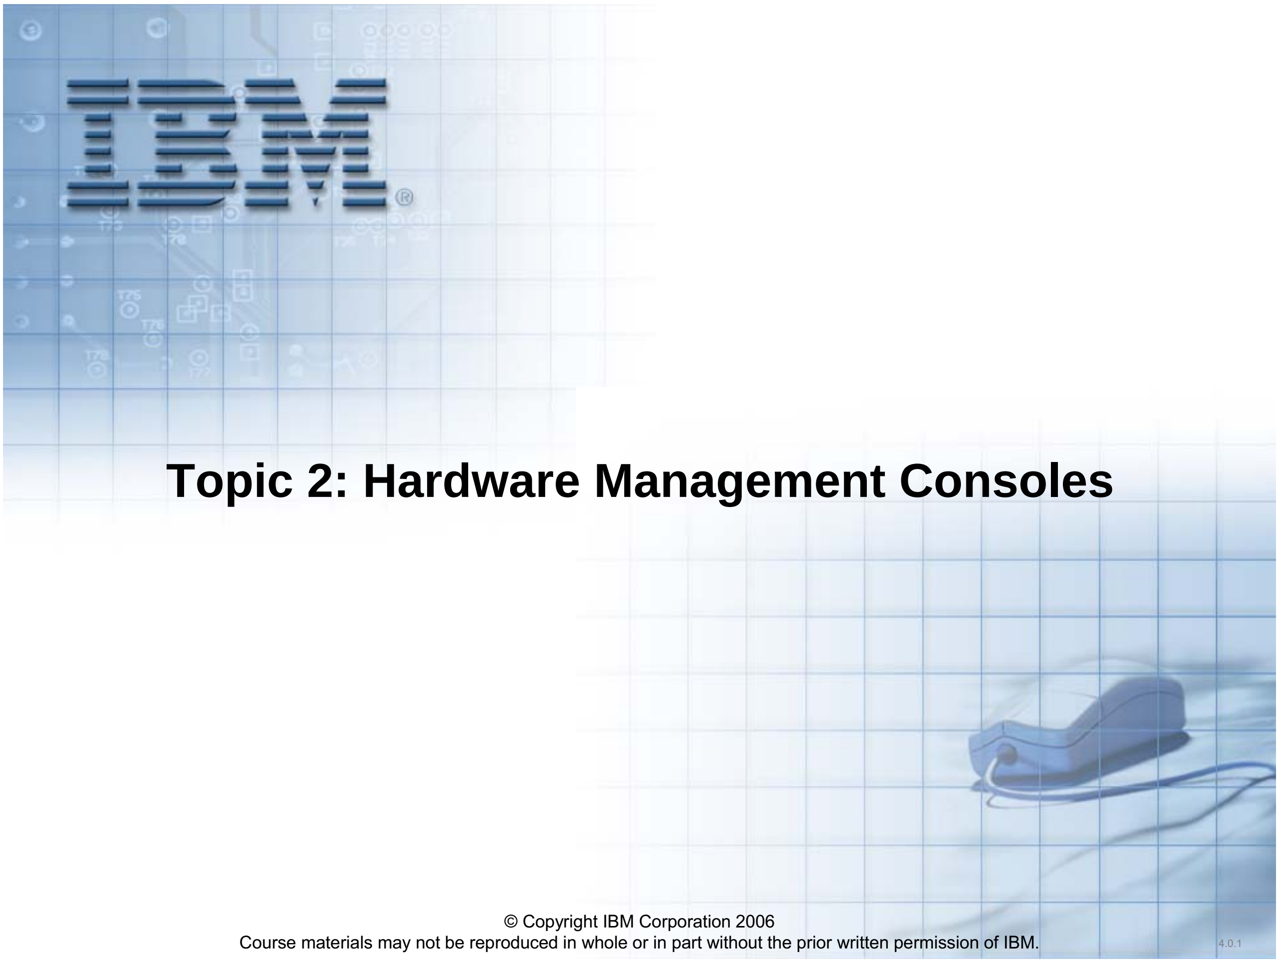

Topic 2: Hardware Management Consoles

4.0.1

Figure 4-27. Topic 2: Hardware Management Consoles AM711.0

Notes:

Virtualization Management Options

- IVM Managed
  - Single server management
  - Entry-level partitioning
    - Limited service function
    - Client LPARs have only virtual devices
    - One profile per partition
  - Limited redundancy
  - Supported on selected models
    - System p5 550Q and below
- HMC managed
  - HMC required
  - Multiple-server management
  - Advanced partitioning
    - Dynamic LPAR, and so forth
    - Physical devices in LPARs
    - Multiple profiles per partition
  - Full redundancy

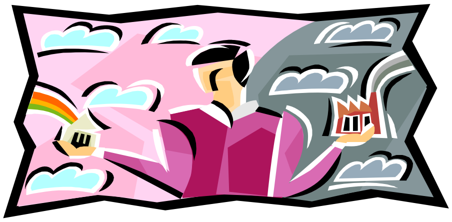

Figure 4-28. Virtualization Management Options AM711.0

Notes:

The HMC offers a service focal point for the systems it controls. It is connected to a dedicated port on the service processor of the POWER5 system using an Ethernet connection. The Ethernet connections from HMC to eServer p5 590 and 595 and eServer i5 595 must be by way of a private LAN, whereas other POWER5 systems can be either private or public. LAN tools are available for determining problems and providing service support, such as call-home and error log notification, through a modem or the Internet.

Multiple POWER5 processor-based servers can be supported by one HMC, located locally or remotely attached.

Hardware Management Console (HMC)

- The HMC helps you manage LPAR configurations.
  - Create and store LPAR profiles that define the processor, memory, and I/O resources allocated to an individual partition
  - Start, stop, and reset a system partition
  - Boot a partition or system by selecting a profile
  - Display system and partition status
  - Display a virtual operator panel of the contents for each partition or controlled system

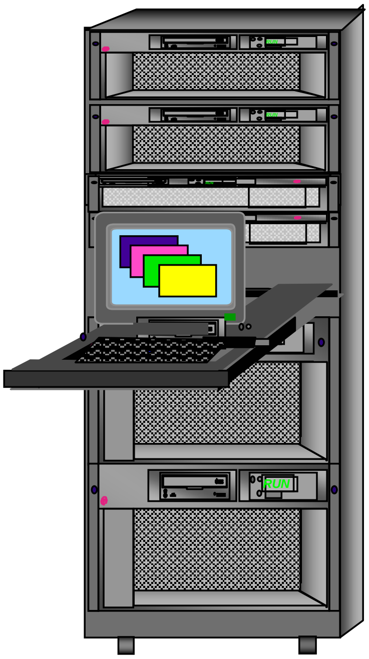

Figure 4-29. Hardware Management Console (HMC) AM711.0

Notes:

The Hardware Management Console (HMC) controls managed systems, including the management of logical partitions and use of Capacity on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information to IBM® for analysis. The HMC provides service technicians with diagnostic information for systems that can operate in a multiple-partitioned environment.

Integrated Virtualization Manager (IVM)

- Available from 10/2005
- Provides HMC functionality in a partition at a much lower price point
- Optional on these IBM systems:
  - IBM System p5 560Q
  - System p5 505, 520, and 550
  - eServer p5 510, 520, and 550
  - OpenPower 710 and 720

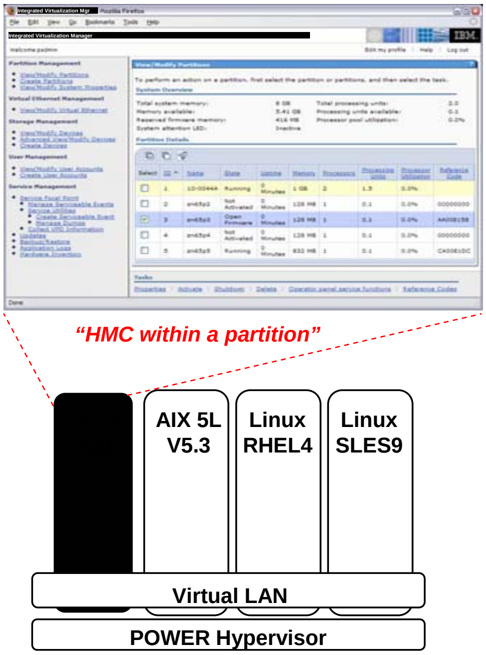

Figure 4-30. Integrated Virtualization Manager (IVM) AM711.0

Notes:

IVM is included at no additional cost with the optional purchase of Advanced POWER virtualization on most IBM System p5™ or IBM eServer™ p5 entry to mid-range models. It is packaged in the Virtual I/O Server.

The Integrated Virtualization Manager is an ideal solution for:

Smaller to mid-sized businesses looking to consolidate existing workloads onto one or two POWER5 servers, or;

Businesses with a geographically distributed IT infrastructure which would otherwise require a Hardware Management Console (HMC) for each location, like a organization with many branch locations).

IVM Limitations

- All physical I/O is owned by the VIOS partition
- Maximum of two VLAN adapters per “client” partition
- One vSCSI adapter per “client” partition
- Limited service function: No concurrent maintenance support
- Uses (some) additional customer compute resources
  - Dedicated boot adapter/device, portion of CPU, memory
  - May not be much greater in relation to base Virtual I/O server

Figure 4-31. IVM Limitations AM711.0

Notes:

Cross Platform Management with IBM Director

- Download IBM Director at no charge*
  - Manage the entire environment from a single point of control
  - Simple point and click GUI plus command line
  - Automation saves time and reduces errors
  - Supports AIX 5L, Linux, and Windows

*https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=dmp

Figure 4-32. Cross Platform Management with IBM Director AM711.0

Notes:

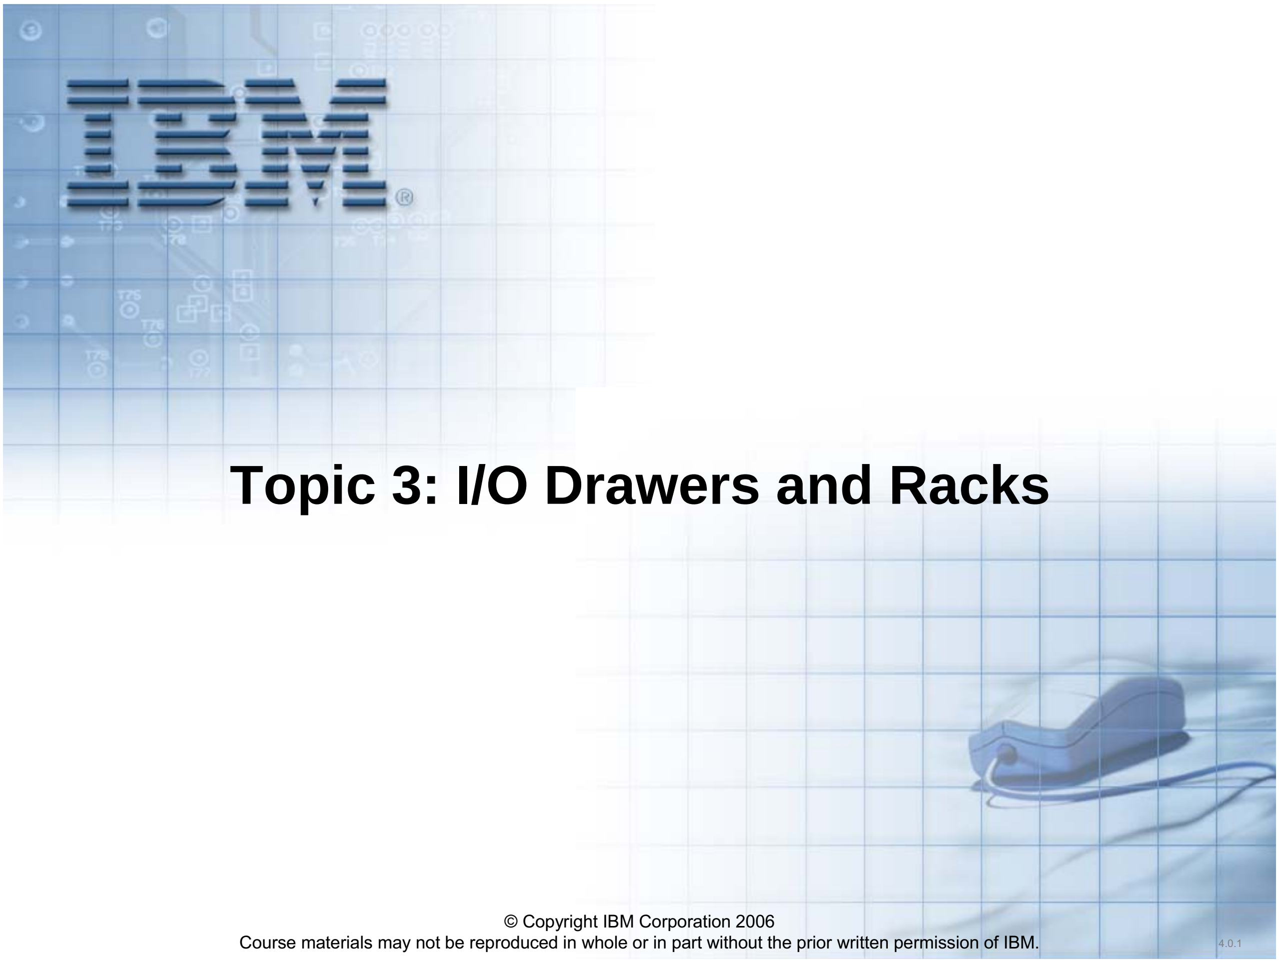

Topic 3: I/O Drawers and Racks

4.0.1

Figure 4-33. Topic 3: I/O Drawers and Racks AM711.0

Notes:

Expandable I/O Subsystems in pSeries

- Provides flexible configuration
- Investment protection
  - Allows for independent upgrades of CEC or I/O
- Minimizes rack footprint and cost

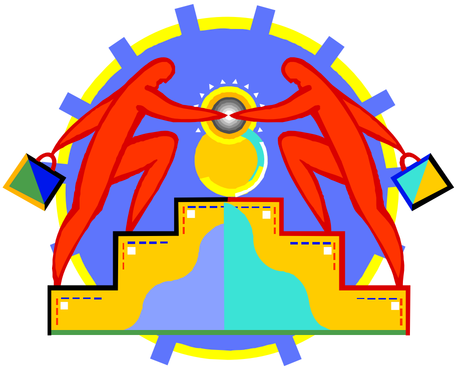

Figure 4-34. Expandable I/O Subsystems in pSeries AM711.0

Notes:

Remote I/O Architecture (RIO2)

- CEC drawer
  - Houses backplane, processors and memory
- I/O drawer
  - Houses PCI-X adapters
- Remote I/O (RIO) cables
  - Connect CEC to I/O drawers

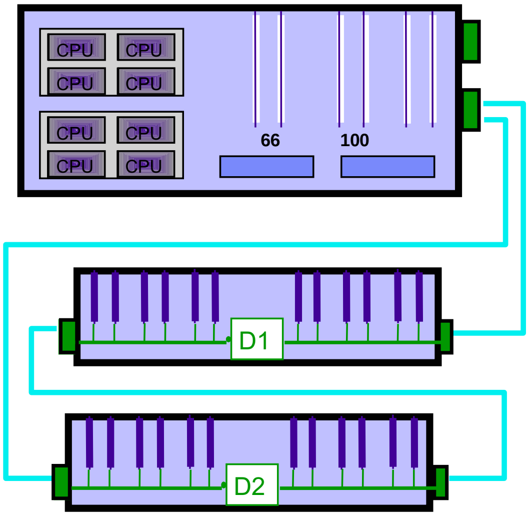

Figure 4-35. Remote I/O Architecture (RIO2) AM711.0

Notes:

I/O Drawer – Model 7311-D20

Front View

4U high

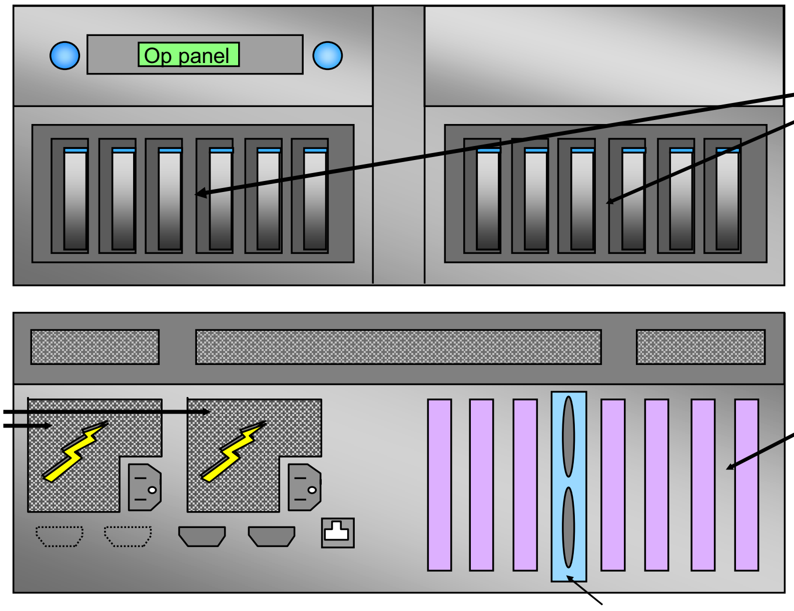

Six-pack hot-swappable disks. 18.2, 36.4, 73.4 or 146.8 GB

Rear View

Two hot-plug power supplies and fans.

Seven PCI-X 64-bit 133 MHz 3.3V hot-plug slots

RIO ports

Figure 4-36. I/O Drawer – Model 7311-D20 AM711.0

Notes:

I/O Drawer – Model 7311-D11

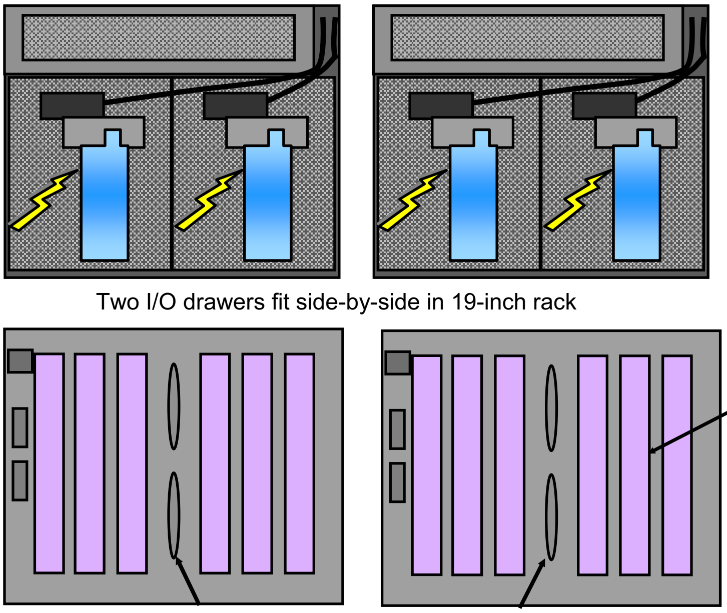

Front View

4U high

Two hot-plug, redundant power supplies for each I/O drawer

Two I/O drawers fit side-by-side in 19-inch rack

Rear View

Six 64-bit 133 MHz @ 3.3v PCI-X hot-plug slots with blind-swap capability.

Two RIO-2 ports per I/O drawer

Figure 4-37. I/O Drawer – Model 7311-D11 AM711.0

Notes:

24-inch I/O Drawer (Features 5791, 5794) 1 of 2

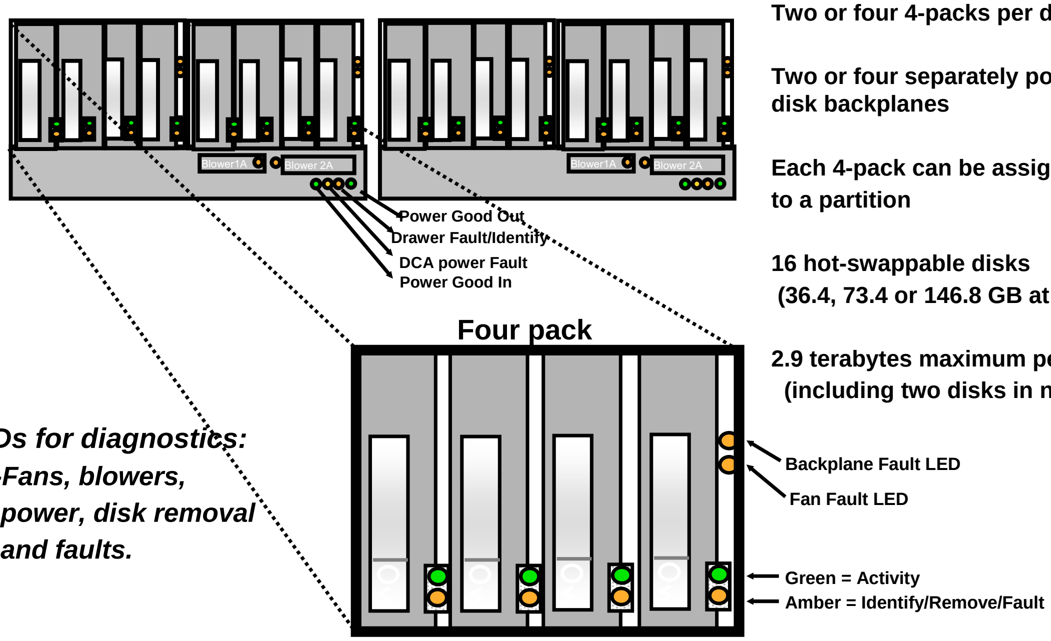

Two or four 4-packs per drawer

Two or four separately powered disk backplanes

Each 4-pack can be assigned to a partition

16 hot-swappable disks (36.4, 73.4 or 146.8 GB at 15 K rpm)

2.9 terabytes maximum per node (including two disks in node)

LEDs for diagnostics:

-Fans, blowers, power, disk removal and faults.

Figure 4-38. 24-inch I/O drawer (features 5791, 5794) 1 of 2 AM711.0

Notes:

24-inch I/O Drawer (Features 5791, 5794) 2 of 2

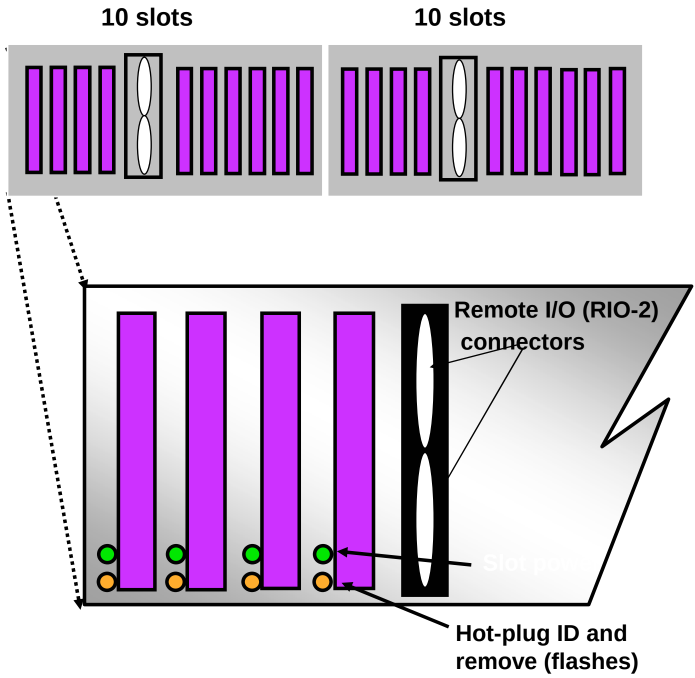
10 slots
10 slots
Remote I/O (RIO-2) connectors
Hot-plug ID and remove (flashes)

- Two x 10 PCI-X slots per drawer, either 32-bit or 64-bit adapters
- Each 10 slots are powered separately
- Each slot may be assigned to a separate partition
- Slots are all hot-plug to allow for application growth or replace faulty adapters while the server remains operational
- Adapter cards are protected by a blind-swap plastic cassette to permit installation without moving drawer and cables

▮ 64-bit PCI-X slots capable of supporting 3.3V signaling PCI-X adapters operating at speeds up to 133 MHz

Figure 4-39. 24-inch I/O drawer (features 5791, 5794) 2 of 2 AM711.0

Notes:

IBM TotalStorage EXP24 (7031-D24, 7031-T24)

- 7031-D24
  - 4U 19” rack model
- 7031-T24
  - Tower model (not shown)
- Up to 24 Ultra320 hot-swappable disk drives

- Disk options:
  - 73.4 / 146.8 / 300 GB 10 K rpm
  - 36.4 / 73.4 / 146.8 GB 15 K rpm
- Four independent disk drive 6-packs
  - Single or dual bus repeater cards for each 6-pack
  - Dual bus repeater cards for high availability
  - Optionally configure as two 12-packs
- Redundant cooling and optional redundant power

Figure 4-40. IBM TotalStorage EXP24 (7031-D24, 7031-T24) AM711.0

Notes:

Server I/O Attachment IBM @server p5 Systems

Server	Drawers	Max.
p5-520	7311-D20	4
p5-550	7311-D20	8
p5-570 Express	7311-D11 7311-D20	12 12
p5-570	7311-D11 7311-D20	20 20
p5-575	F/C 5791 internal drawer F/C 5794 internal drawer 7040-61D*	1 1 1
p5-590	F/C 5791 internal drawer F/C 5794 internal drawer 7040-61D*	8 8 8
p5-595	F/C 5791 internal drawer F/C 5794 internal drawer 7040-61D*	12 12 12

Figure 4-41. Server I/O Attachment IBM @server p5 Systems AM711.0

Notes:

*7040-61D may be attached if migrated from another system but not orderable as new.

Server I/O Attachment IBM System p5

Server	Drawers	Max.
p505 Express **p5 510/510Q**	**none**	**none**
p5 520 Express	**7311-D20**	**4**
p5 520Q Express	**7311-D20**	**4**
p5 550 Express	**7311-D20**	**8**
p5 550Q Express	**7311-D20**	**8**
p5 560Q	**none**	**none**
p5 570	**7311-D11,7311-D20**	**20**

Figure 4-42. Server I/O Attachment IBM System p5 AM711.0

Notes:

Information from Facts and Features p. 13

7040-61D may be attached if migrated from another system but not orderable as new.

IBM System p5 Racks

For rack-mount
installation of p5-520/550/550Q

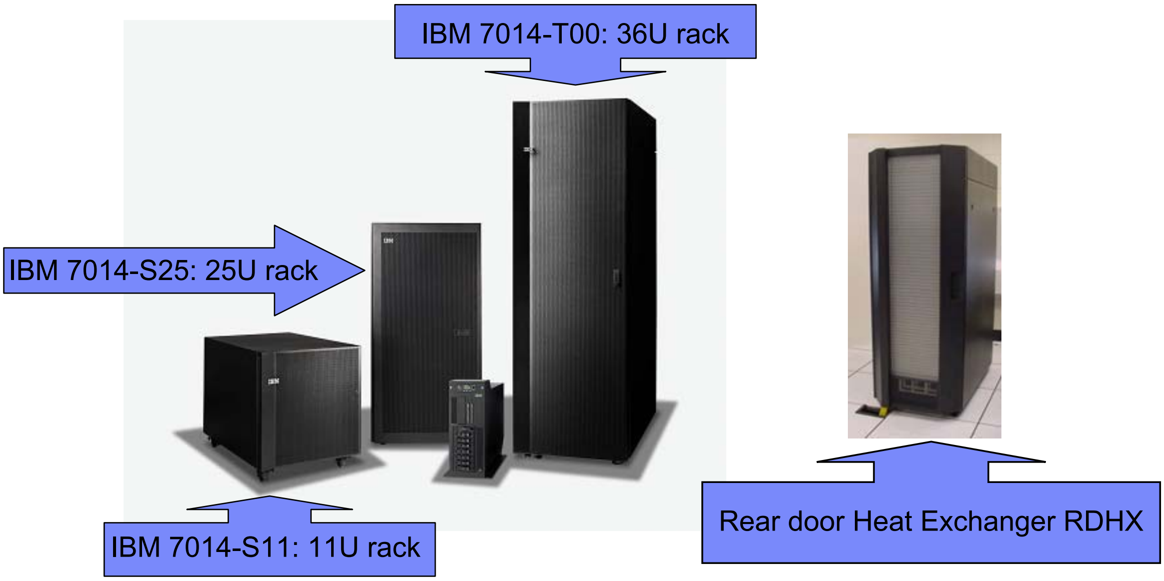

IBM 7014-T00: 36U rack

IBM 7014-S25: 25U rack

IBM 7014-S11: 11U rack

Rear door Heat Exchanger RDHX

Figure 4-43. IBM System p5 racks AM711.0

Notes:

These are the choices for racks now… with support for p5-520/550/550Q

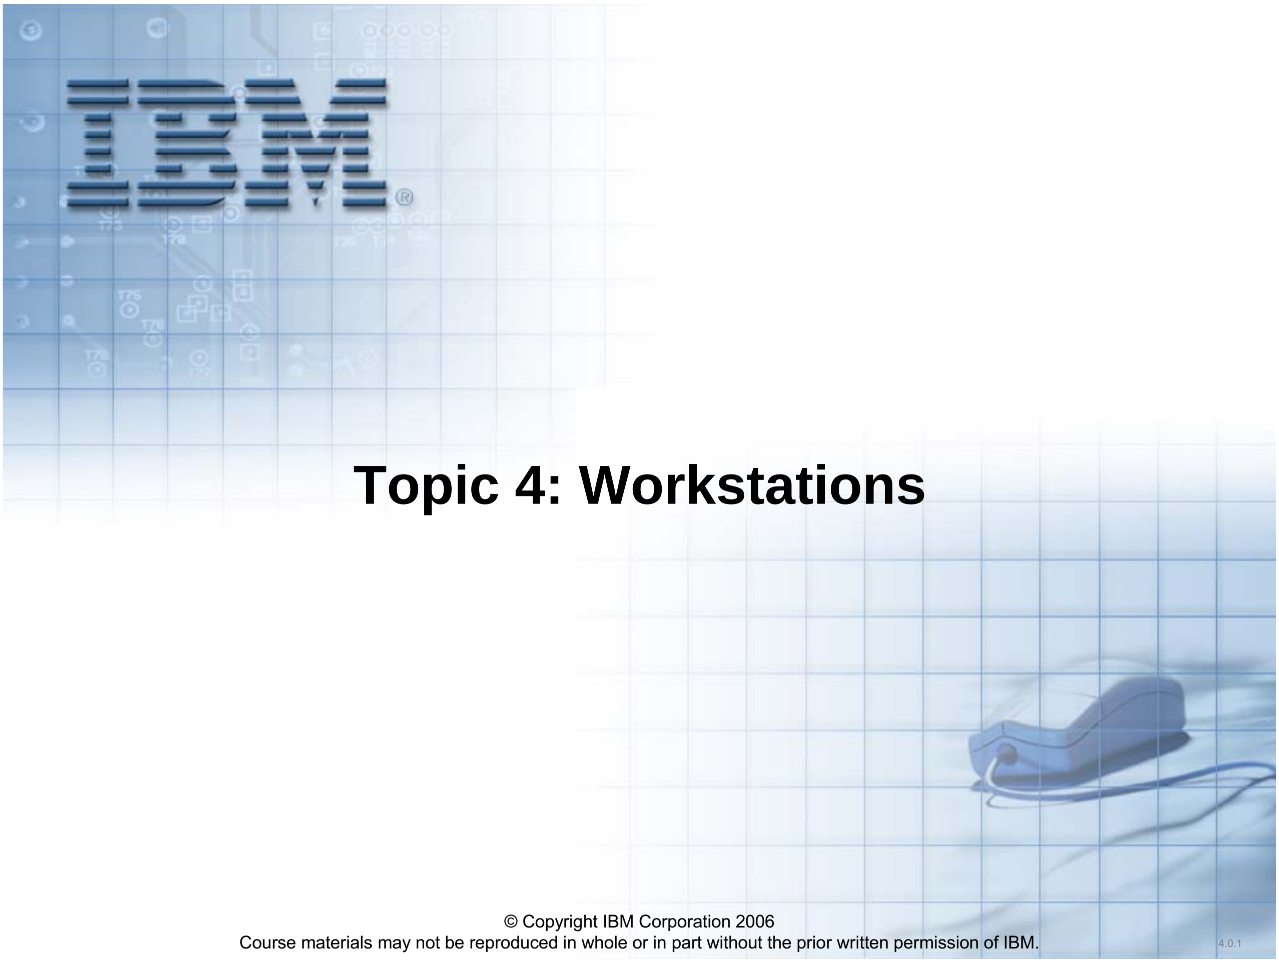

Topic 4: Workstations

4.0.1

Figure 4-44. Topic 4: Workstations AM711.0

Notes:

IBM IntelliStation POWER 185

IBM PowerPC 970 GRAPHICS WORKSTATION

- One-core 2.5 GHz
- Two-core 2.5 GHz
- New, quieter packaging

VERY FOCUSED TARGET ENVIRONMENT

- CATIA V4 and V5 install base
- Automotive and Aerospace
- High Performance Computing

KEY BUYING CRITERIA

- Low price with good performance
- Quiet
- 64-bit CATIA V5 support

DELIVER BUSINESS VALUE

- Increased design productivity / creativity
- Supports large models and assemblies
- AIX 5L platform designed for reliability and security
- System administration consistent with servers

IBM's entry POWER workstation

Figure 4-45. IBM IntelliStation POWER 185 AM711.0

Notes:

IBM IntelliStation Power 285

- Exceptional performance for high-end design and analysis applications
- Quiet operation allows users to focus on their work in comfort
- Designed for 64-bit CATIA V5 to solve complex engineering problems
- Variety of memory and storage options for demanding engineering environments

NEW!

July 2006

One or two POWER5+ processors running at 1.9 GHz or 2.1 GHz

Figure 4-46. IBM IntelliStation Power 285 AM711.0

Notes:

One or two 64-bit POWER5+ microprocessors running at 1.9 GHz or 2.1 GHz

Quiet deskside

2D graphics (GXT135P)

3D graphics (GXT4500P and GXT6500P)

3D Input (SpaceBall Plus 3D, SpaceMouse Plus 3D, SpacePilot)

15”, 17”, 19”, and 20” TFT Displays

Up to 32GB memory

6 PCI-X slots

2 USB ports, 2 serial ports

HMC service processor

Very quiet (5.0 bels)

AIX 5L™ V5.2 and AIX 5L V5.3

Graphics Accelerators

- POWER GXT135P
  - Versatile, low-priced 2D graphics accelerator
- GXT4500P
  - Value-oriented price/performance for entry 3D workstations.
- GXT6500P
  - The highest performing graphics accelerator
  - Supports the higher performance requirements of the more demanding user of advanced 3D workstations

Figure 4-47. Graphics Accelerators AM711.0

Notes:

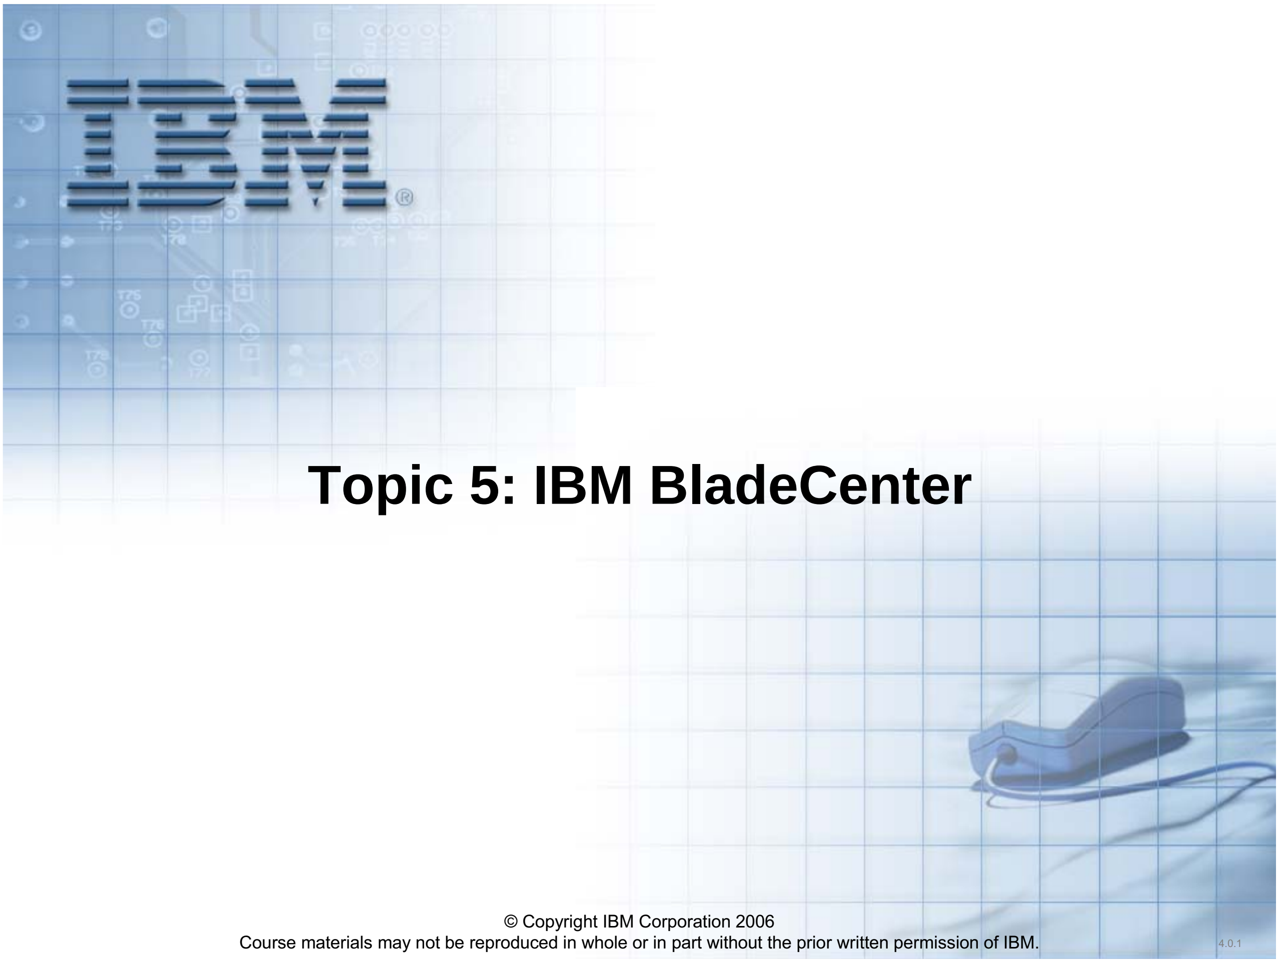

Topic 5: IBM BladeCenter

Figure 4-48. Topic 5: IBM BladeCenter AM711.0

Notes:

Selling IBM BladeCenter

- IBM has a technology advantage
  - Lead over competitors' products
  - Great references
- It fits our clients business strategies
  - Server consolidation
  - Infrastructure simplification
- The world is moving towards blades
  - BladeCenter sales are expected to outperform other products.

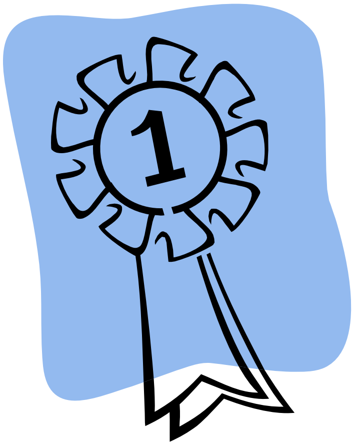

Figure 4-49. Selling IBM BladeCenter AM711.0

Notes:

IBM gained share in the blade server space YTY and retained its leadership position for the 7th straight quarter.

Optimized Workloads for BladeCenter

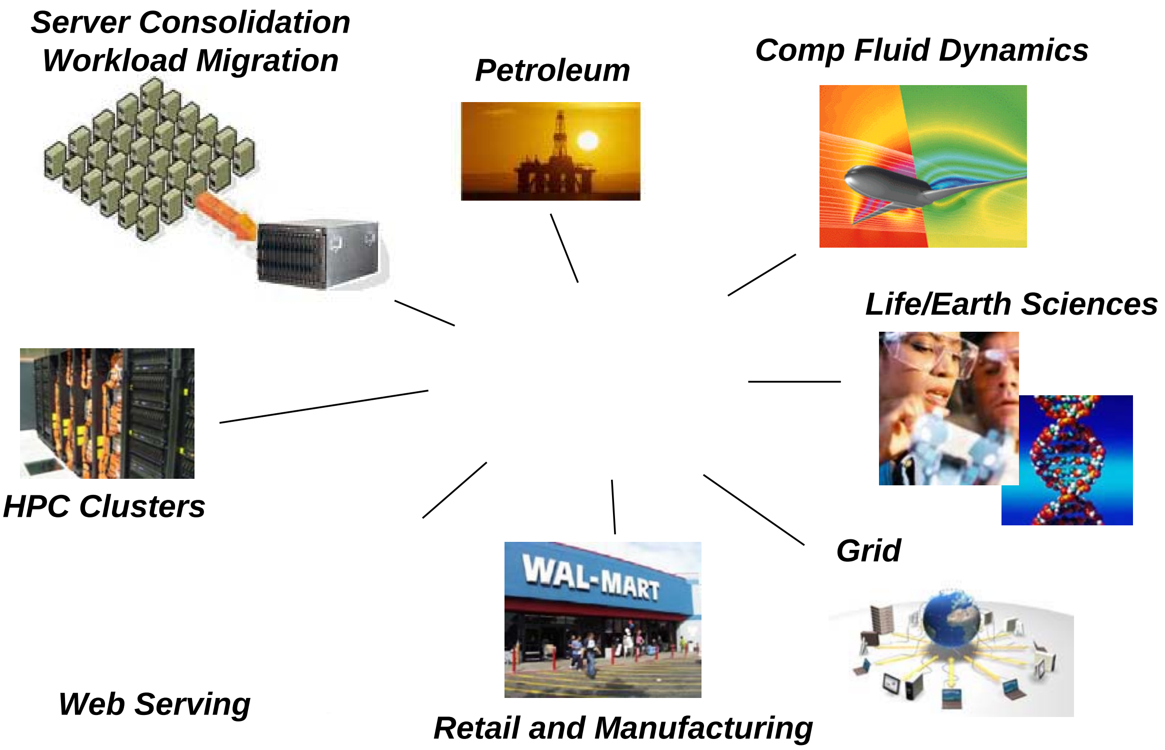

Server Consolidation
Workload Migration

Petroleum

Comp Fluid Dynamics

Life/Earth Sciences

HPC Clusters

Grid

Web Serving

Retail and Manufacturing

Figure 4-50. Optimized Workloads for BladeCenter AM711.0

Notes:

IBM BladeCenter

- 7U Chassis
  - Support for 14 blades
  - Up to 168 blades in 42U rack
  - Hot-plug components with full redundancy
    - Passive mid-plane
    - Power, cooling
    - Ethernet and Fibre Channel modules
    - Management modules
  - Telco version available (BladeCenter T)

- 9U chassis
- All the functionality of BladeCenter plus:
  - 4X InfiniBand-capable fabrics
  - Larger power supplies
  - More robust thermal
- Maintains compatibility with current blades and switches.

BladeCenter H

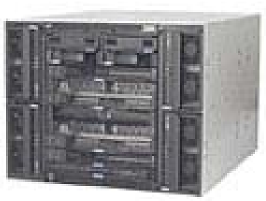

BladeCenter T

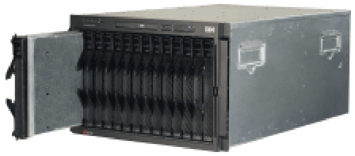

BladeCenter

Figure 4-51. IBM BladeCenter AM711.0

Notes:

IBM BladeCenter JS20

- BladeCenter JS20 Telco Model 8842-RTX
  - 2 2.2 GHz PowerPC 970FX cores
  - 512 KB L2 cache
  - 2 GB – 8 GB DDR SDRAM memory
  - NEBS 3 compliant when installed in BladeCenter T chassis
- Prices of AIX5.2 and AIX5.3 reduced to $85/core for processor category C5 (JS20 and JS21)
- JS20 Models 8842-42X and 8842-4TX EOM June 30, 2006

Figure 4-52. IBM BladeCenter JS20 AM711.0

Notes:

IBM BladeCenter JS21 and JS21 Express

- Delivers up to 3x the performance of the JS20*
- Delivers up to twice the performance of HP BL60p Itanium blade*
- First BladeCenter blade server with built-in virtualization (APV)**
- First blade server designed for 10 Gb-capable BladeCenter H
- Greater reliability and performance with SAS hard drives
- DDR2 memory
- Integrated PCI-Express

*SPECint_rate2000: JS20 (21) vs. JS21 (67.9) 1 SPECint_rate2000: HP BL60p (33.4) vs. JS21 (67.9) Source: www.spec.org IBM SPEC results submitted to SPEC on 2/09/06 Competitive SPEC results current as of 2/08/06 2 Must acquire optional VIOS license to utilize Micro-partitioning

Figure 4-53. IBM BladeCenter JS21 and JS21 Express AM711.0

Notes:

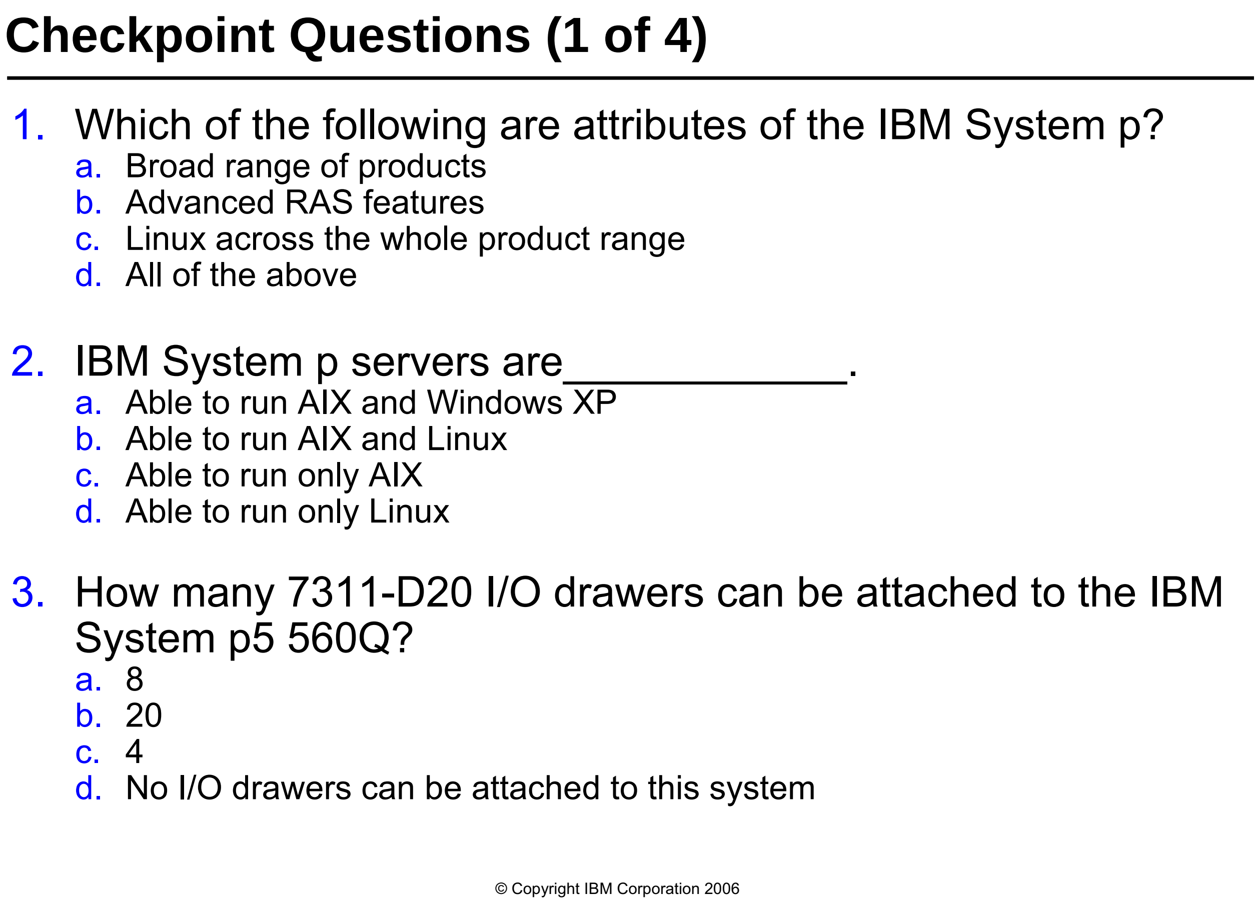

Checkpoint Questions (1 of 4)

1. Which of the following are attributes of the IBM System p?
   a. Broad range of products
   b. Advanced RAS features
   c. Linux across the whole product range
   d. All of the above

2. IBM System p servers are____________.
   a. Able to run AIX and Windows XP
   b. Able to run AIX and Linux
   c. Able to run only AIX
   d. Able to run only Linux

3. How many 7311-D20 I/O drawers can be attached to the IBM System p5 560Q?
   a. 8
   b. 20
   c. 4
   d. No I/O drawers can be attached to this system

Figure 4-54. Checkpoint Questions (1 of 4) AM711.0

Notes:

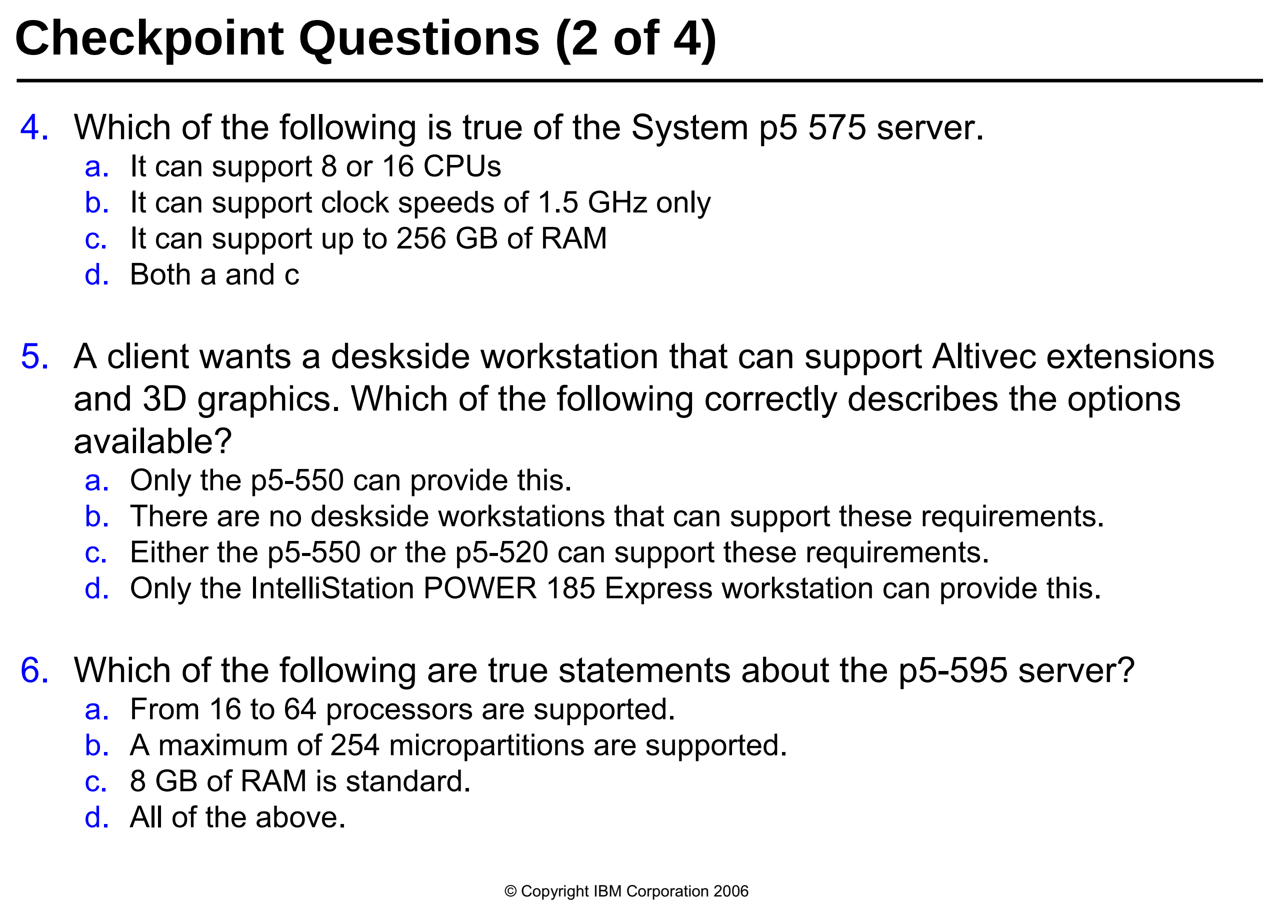

Checkpoint Questions (2 of 4)

4. Which of the following is true of the System p5 575 server.
   a. It can support 8 or 16 CPUs
   b. It can support clock speeds of 1.5 GHz only
   c. It can support up to 256 GB of RAM
   d. Both a and c

5. A client wants a deskside workstation that can support Altivec extensions and 3D graphics. Which of the following correctly describes the options available?
   a. Only the p5-550 can provide this.
   b. There are no deskside workstations that can support these requirements.
   c. Either the p5-550 or the p5-520 can support these requirements.
   d. Only the IntelliStation POWER 185 Express workstation can provide this.

6. Which of the following are true statements about the p5-595 server?
   a. From 16 to 64 processors are supported.
   b. A maximum of 254 micropartitions are supported.
   c. 8 GB of RAM is standard.
   d. All of the above.

Figure 4-55. Checkpoint Questions (2 of 4) AM711.0

Notes:

Checkpoint Questions (3 of 4)

7. Which of the following statements is true about the IBM System p5 570 server?
    a. It can support processor speeds of 1.9 GHz and 2.2 GHz
    b. It can support up to 16 processor cores per system
    c. It can support up to 20 I/O drawers
    d. All of the above.
8. Which server offers quad-core technology in a 2U package?
    a. p5-510Q
    b. p5-550Q
    c. p5-570Q
    d. No current system provides this capability.

Figure 4-56. Checkpoint Questions (3 of 4) AM711.0

Notes:

Checkpoint Questions (4 of 4)

9. Which of the following statements is true about the JS20 Blade.
   a. It can run both AIX and Linux.
   b. It can support micropartitions.
   c. It can support SCSI disks internally.
   d. All of the above.

10. The D20 I/O drawer provides support for:
   a. 12 hotplug SCSI disks and 7 PCI-X slots
   b. 6 hotplug SCSI disks and 6 PCI-X slots
   c. 14 PCI-X slots
   d. None of the above

Figure 4-57. Checkpoint Questions (4 of 4) AM711.0

Notes:

Unit Summary

Having completed this unit, you should be able to:

- Describe the advantages to the client of the IBM System p servers and related products.

- Identify and position the current range of:
  - IBM System p5 servers
  - IBM IBM BladeCenter servers
  - IBM IntelliStation POWER workstations
  - IBM System p I/O drawers and racks

- Identify and position the HMC and IVM methods of managing partitioned servers.

Figure 4-58. Unit Summary AM711.0

Notes:

Any Questions?

Figure 4-59. Any Questions? AM711.0

Notes:

Unit 5. Infrastructure

What This Unit Is About

This unit describes the network, cluster and grid possibilities associated with System p solution design.

What You Should Be Able to Do

After completing this unit, you should be able to:

- List the main features of the System p network support
- List the network adapters supported by System p
- Identify appropriate adapters from customer requirements

- Describe what can be meant by the term cluster
- Describe the benefits of clusters
- List the components of the Cluster 1600 solution
- Describe the current Cluster 1600 solution offerings

- Describe what Grid computing is
- Know why and when to sell Grid computing

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

After completing this unit, you should be able to:

- List the main features of the System p network support.
- List the network adapters supported by System p.
- Identify appropriate adapters from customer requirements.

- Describe what can be meant by the term cluster.
- Describe the benefits of clusters.
- List the components of the Cluster 1600 solution.
- Describe the current Cluster 1600 solution offerings.

- Describe what Grid computing is.
- Know why and when to sell Grid computing.

Figure 5-1. Unit Objectives AM711.0

Notes:

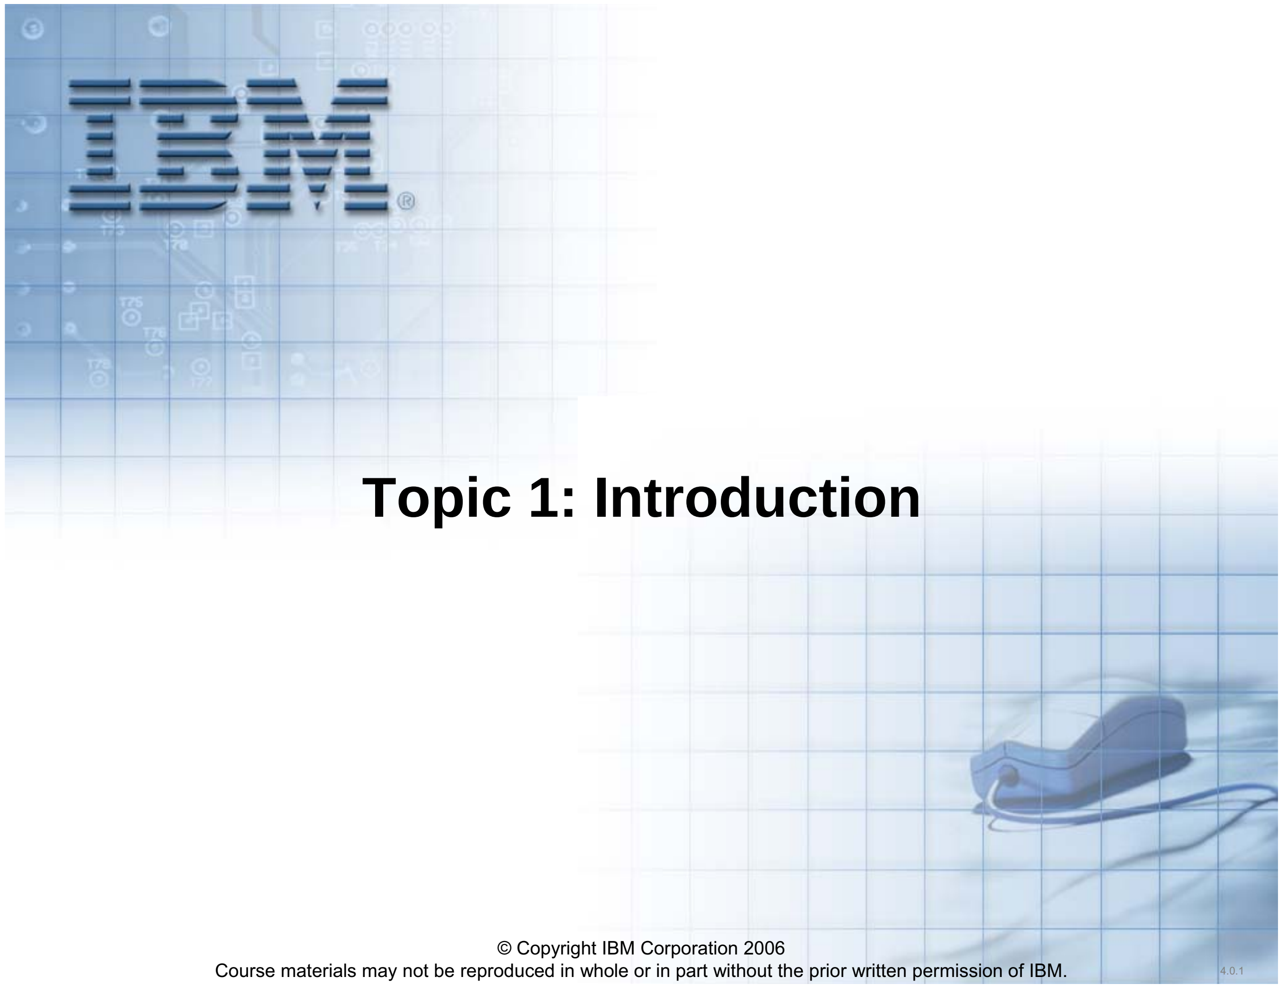

Topic 1: Introduction

Figure 5-2. Topic 1: Introduction AM711.0

Notes:

Introduction

- In this topic we discuss System p products which support our clients' requirements for infrastructure. This is divided into three topics:

Networks

Clusters

Grids

Figure 5-3. Introduction AM711.0

Notes:

What Do We Mean By Infrastructure?

- The network attachments
- The storage devices*
- The cluster interconnects
- The Grid components

* Storage is a big subject in its own right and is covered in its own topic.

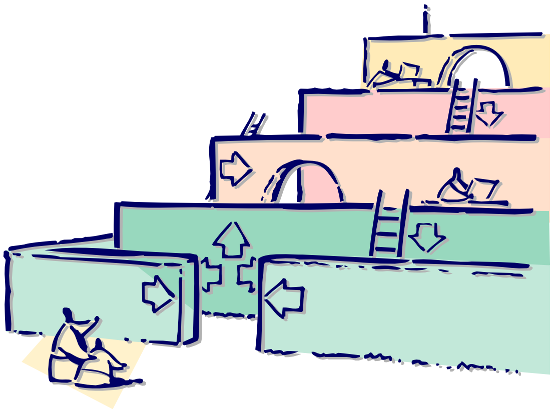

Figure 5-4. What Do We Mean By Infrastructure? AM711.0

Notes:

For this discussion we will consider infrastructure as the ‘components' that allow clients to access their pSeries systems.

*Storage is a big subject in its own right and is covered in its own topic.

The Value of Infrastructure

- Networks provide the means for users to connect to applications.
- Work can be distributed among groups of computers as in clusters or Grids.
- Resilience can be increased for applications by creating failover clusters.
- Groups of computers can be managed from a single location.

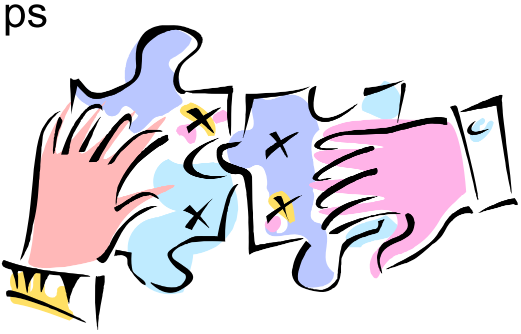

Figure 5-5. The Value of Infrastructure AM711.0

Notes:

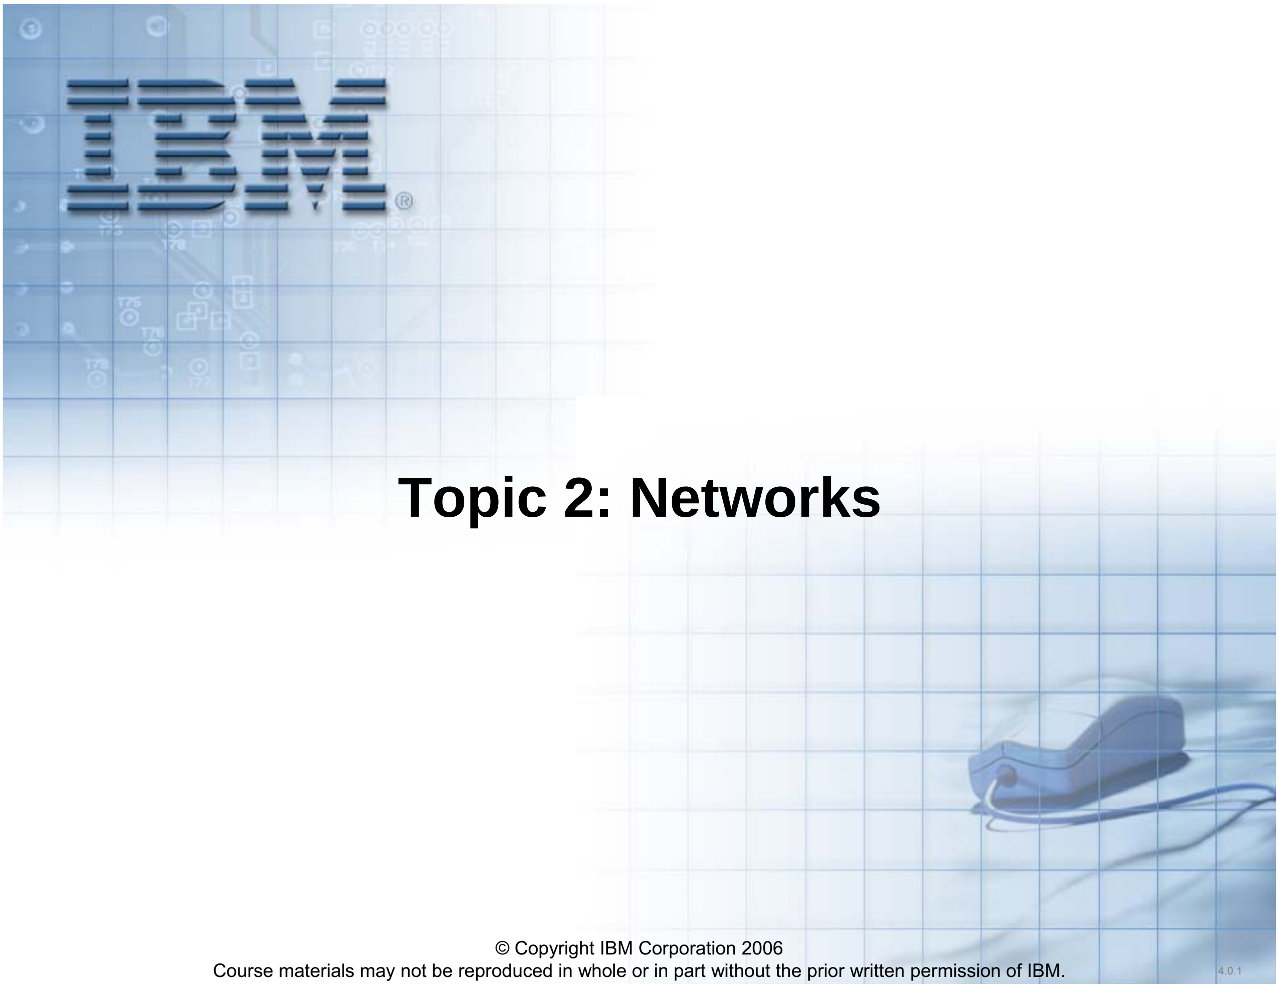

Topic 2: Networks

Figure 5-6. Topic 2: Networks AM711.0

Notes:

Topic Objectives

- List the benefits of networking.
- List the main features of System p network support.
- List the network adapters supported by System p.
- Identify appropriate adapters from customer requirements.

Figure 5-7. Topic Objectives AM711.0

Notes:

Networks Are Everywhere!

Networked computers allow:

- Access to the Internet
- Shared access to resources
  - Printers
  - Disk storage
  - Tape storage
- Collaborative computing
  - E-mail
  - Calendar
  - Net meetings
  - Team rooms
- Simplification of the connection between different
  - Platforms
  - Operating systems
  - Network protocols

Figure 5-8. Networks Are Everywhere! AM711.0

Notes:

Advantages of System p Network Products

☑Easy to integrate into existing network infrastructure

☑Support for many adapter types

☑Support for a variety of network protocols

☑Advanced network features of operating systems

Figure 5-9. Advantages of System p Network Products AM711.0

Notes:

Buzzwords and Acronyms

- Networks are described in various ways:
  - Topology
    - LAN
    - MAN
    - WAN
  - Network Type
    - Ethernet
    - Token-ring
    - ATM
  - Protocol
    - TCP/IP
    - X.25
  - Physical connection
    - Copper
    - Fibre-optic

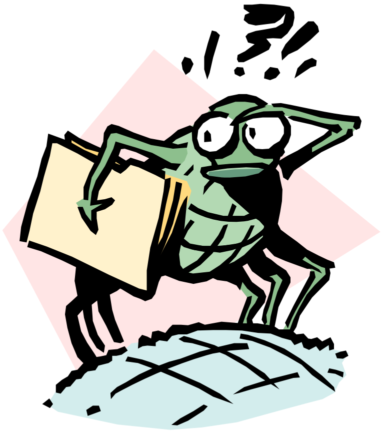

Figure 5-10. Buzzwords and Acronyms AM711.0

Notes:

Copper versus Fiber Optic

- Pros of copper
  - Currently the most cost effective option
  - Customers can migrate to Gigabit Ethernet using existing wiring
- Pros of fiber
  - Longer distances
  - Impervious to electromagnetic interference
  - Signal security – immune to electronic tapping
  - Greater bandwidth – 10 Gbps versus 1 Gbps
  - Lighter in weight, thinner, and more durable than copper

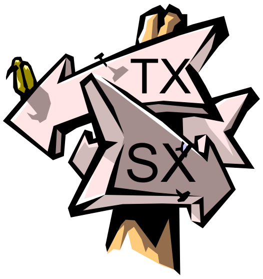

Figure 5-11. Copper versus Fiber Optic AM711.0

Notes:

Adapters which connect to copper cables are generally designated as –TX prefix (UTP).

Copper cable is limited to distance of 100 m.

Adapters which connect to fiber optic cable are designated:

SX prefix. Multimode Fiber optic cable. Max. distance 550 m

LX prefix. Single mode Fiber optic cable. Max. distance 5 Km

Network Adapters

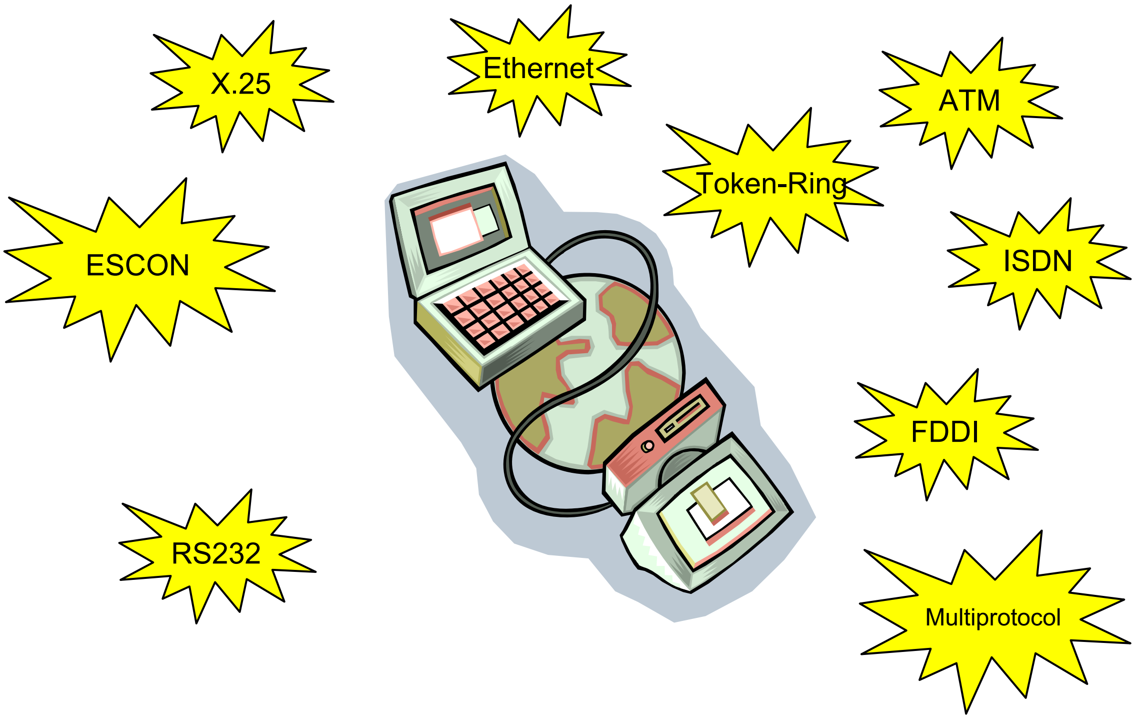

X.25

Ethernet

ATM

Token-Ring

ESCON

ISDN

FDDI

RS232

Multiprotocol

Figure 5-12. Network Adapters AM711.0

Notes:

Ethernet Adapter Choices

- Single port
- Multi-port
- Integrated
- PCI-X
- Fiber
- Copper

Figure 5-13. Ethernet Adapter Choices AM711.0

Notes:

Integrated Ethernet Adapters

10/100/1000 Mbps two ports
- System p5 185 Express
- System p5 505 Express
- System p5 510/510Q Express
- System p5 520/520Q Express
- System p5 550/550Q Express
- System p5 560Q Express

10/100/1000 Mbps four ports
- System p5 560Q*
- System p5 575

10/100/1000 Mbps eight ports
- System p5 570*

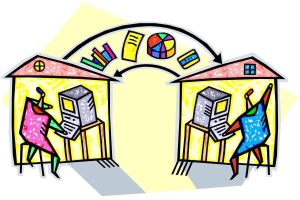

*Assuming maximum building blocks installed

Figure 5-14. Integrated Ethernet Adapters

Notes:

Single-port PCI-X Ethernet Adapters

- 10/100 Mbps Ethernet PCI II
- 10/100/1000 Mbps Base-TX Ethernet PCI-X
- Gigabit Ethernet-SX PCI-X
- 10 Gigabit Ethernet –PCI-X
- 10 Gigabit Ethernet-LR –PCI-X

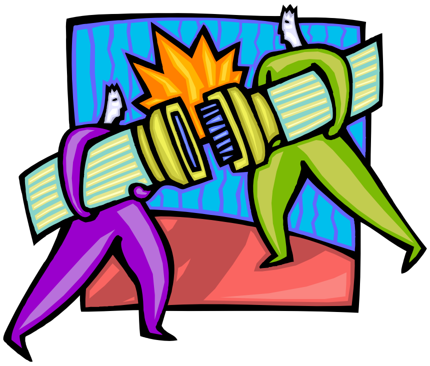

Figure 5-15. Single-port PCI-X Ethernet Adapters

AM711.0

Notes:

Multiport PCI-X Ethernet Adapters

- Requirement
  - Minimize the number of PCI slots used for network connections

- Solution
  - 4-port 10/100 Mbps Ethernet
  - 2-port 10/100/1000 Mbps Base-TX Ethernet PCI-X
  - 2-port Gigabit Ethernet-SX PCI-X

  - 4-port 10/100/1000 Mbps PCI-X Ethernet adapter
    - Two integrated dual-port controllers on a single adapter

New

Figure 5-16. Multiport PCI-X Ethernet Adapters AM711.0

Notes:

Asynchronous Transfer Mode (ATM)

- Requirement:
  - High speed
  - Error-resistant
  - Applications such as:
    - Streaming audio
    - Streaming video

- Solution:
  - 64-bit ATM 155 Mbps (UTP, or Fiber)
  - ATM 622 Mbps (Fiber)

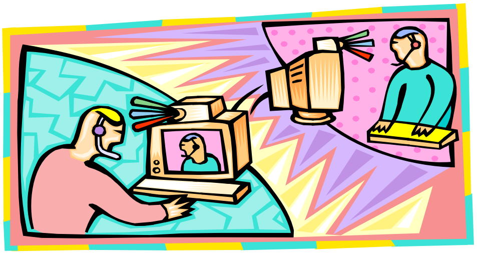

Figure 5-17. Asynchronous Transfer Mode (ATM) AM711.0

Notes:

Other Network Connections

- Requirement
  - Connection to existing token-ring network environment
- Solution:
  - Token-Ring Adapter 4/16 Mbps
- Requirement:
  - Connection to S/390 mainframe
- Solution:
  - ESCON control unit

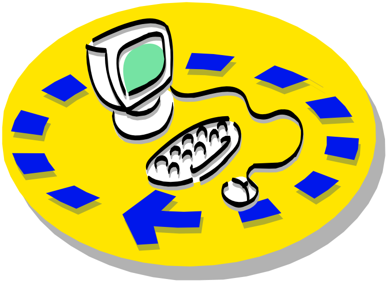

Figure 5-18. Other Network Connections

AM711.0

Notes:

Other Network Adapters

- Requirement:
  - Cluster 1600 connection
- Solution:
  - InfiniBand switch
  - High-performance switch Network interface
  - SP Switch2 system attachment

- Requirement:
  - Voice response systems
- Solution:
  - Quad digital trunk telephony

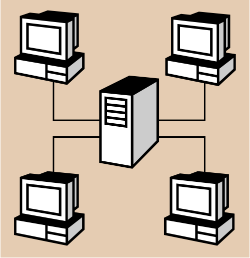

Figure 5-19. Other Network Adapters AM711.0

Notes:

Other Network Adapters

- Requirement:
  - Connection to serial terminals
- Solution:
  - 8/128 port asynchronous controller

- Requirement:
  - Secure IPSec virtual private network
- Solution:
  - PCI cryptographic coprocessor

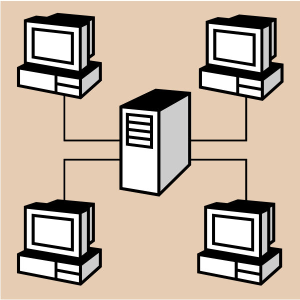

Figure 5-20. Other Network Adapters AM711.0

Notes:

Virtual Ethernet

- No physical adapter required
- Inter-partition communication support
- IP forwarding provided by I/O server partition
- Can have multiple connections per partition

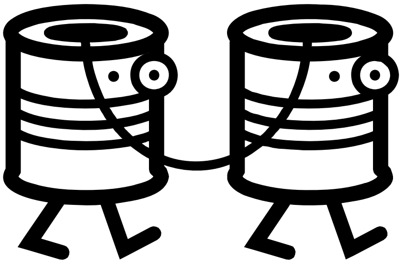

Figure 5-21. Virtual Ethernet AM711.0

Notes:

Operating System Support

- Support for all the adapters covered previously

- Support for the most popular Ethernet adapters
- For other adapters check with the distribution of Linux and version

Figure 5-22. Operating System Support

AM711.0

Notes:

Checkpoint Questions (1 of 6)

1. Which of the following is not a benefit of networking?
   a. Shared access to resources
   b. Collaborative computing
   c. Simplification of connection between technologies
   d. Increased system security

2. Which of the following is not true about System p?
   a. Support for many network adapter types
   b. Support only for IBM proprietary adapters
   c. Operating systems support for advanced networking features
   d. Ease of integration into existing network infrastructure

3. True or False. A single copper cable segment can carry Ethernet traffic over 20 kilometers.

4. True or False. TCP/IP is a suite of protocols.

Figure 5-23. Checkpoint Questions (1 of 6) AM711.0

Notes:

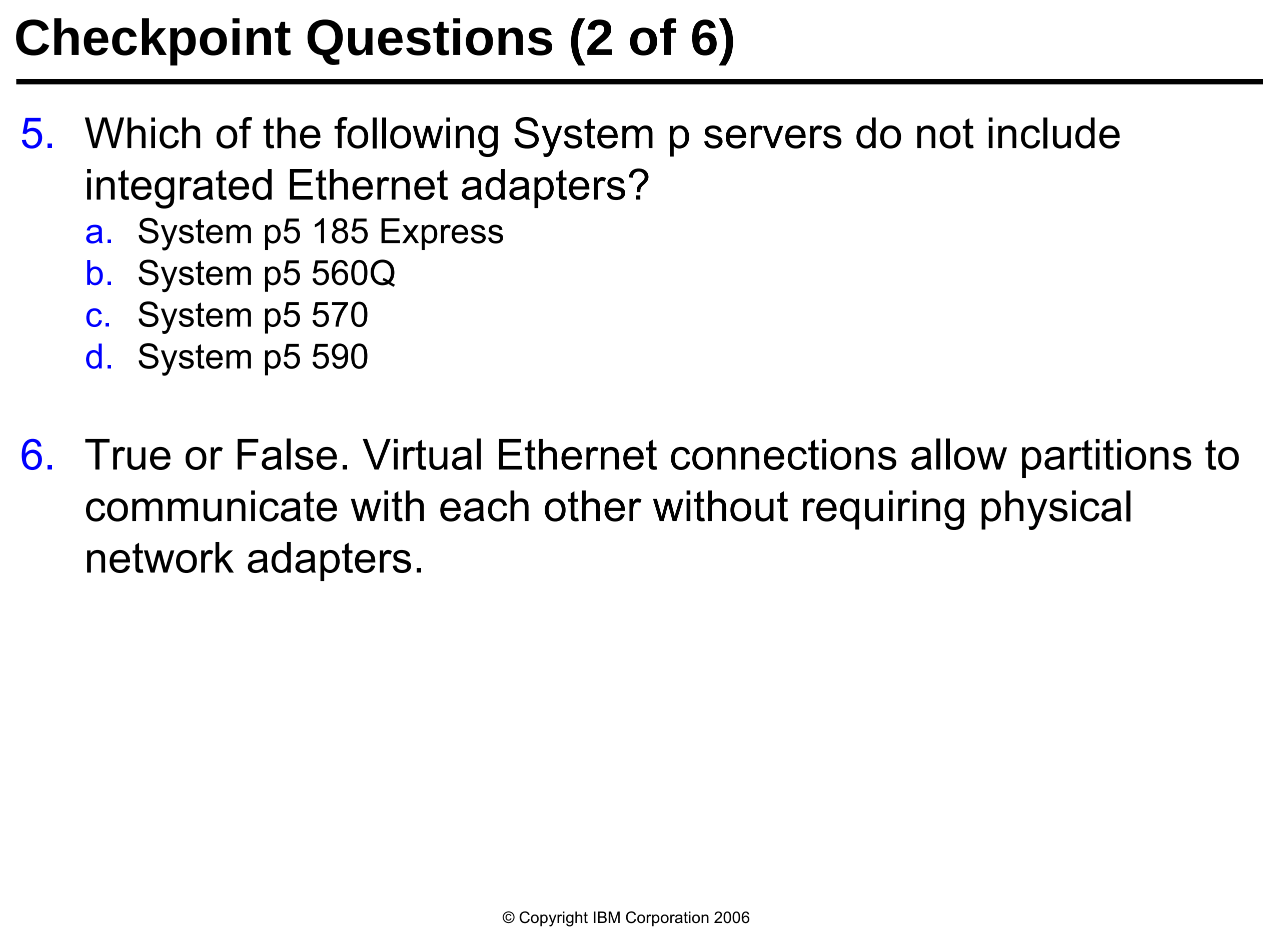

Checkpoint Questions (2 of 6)

5. Which of the following System p servers do not include integrated Ethernet adapters?
   a. System p5 185 Express
   b. System p5 560Q
   c. System p5 570
   d. System p5 590

6. True or False. Virtual Ethernet connections allow partitions to communicate with each other without requiring physical network adapters.

Figure 5-24. Checkpoint Questions (2 of 6) AM711.0

Notes:

Topic Summary

You should now be able to:

- List the benefits of networking.
- List the main features of System p network support.
- List the network adapters supported by System p.
- Identify appropriate adapters from customer requirements.

Figure 5-25. Topic Summary AM711.0

Notes:

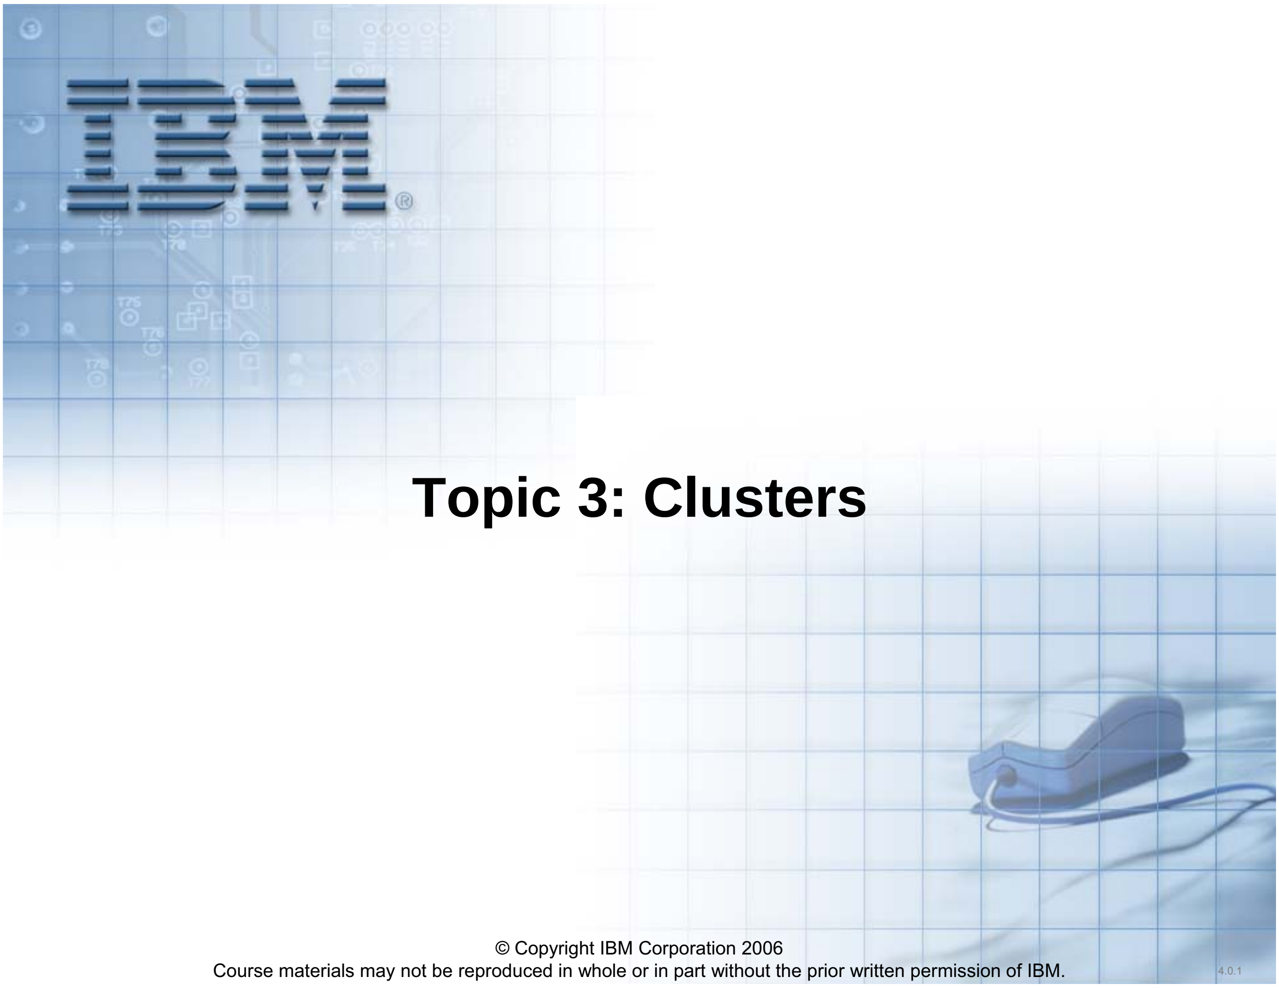

Topic 3: Clusters

4.0.1

Figure 5-26. Topic 3: Clusters AM711.0

Notes:

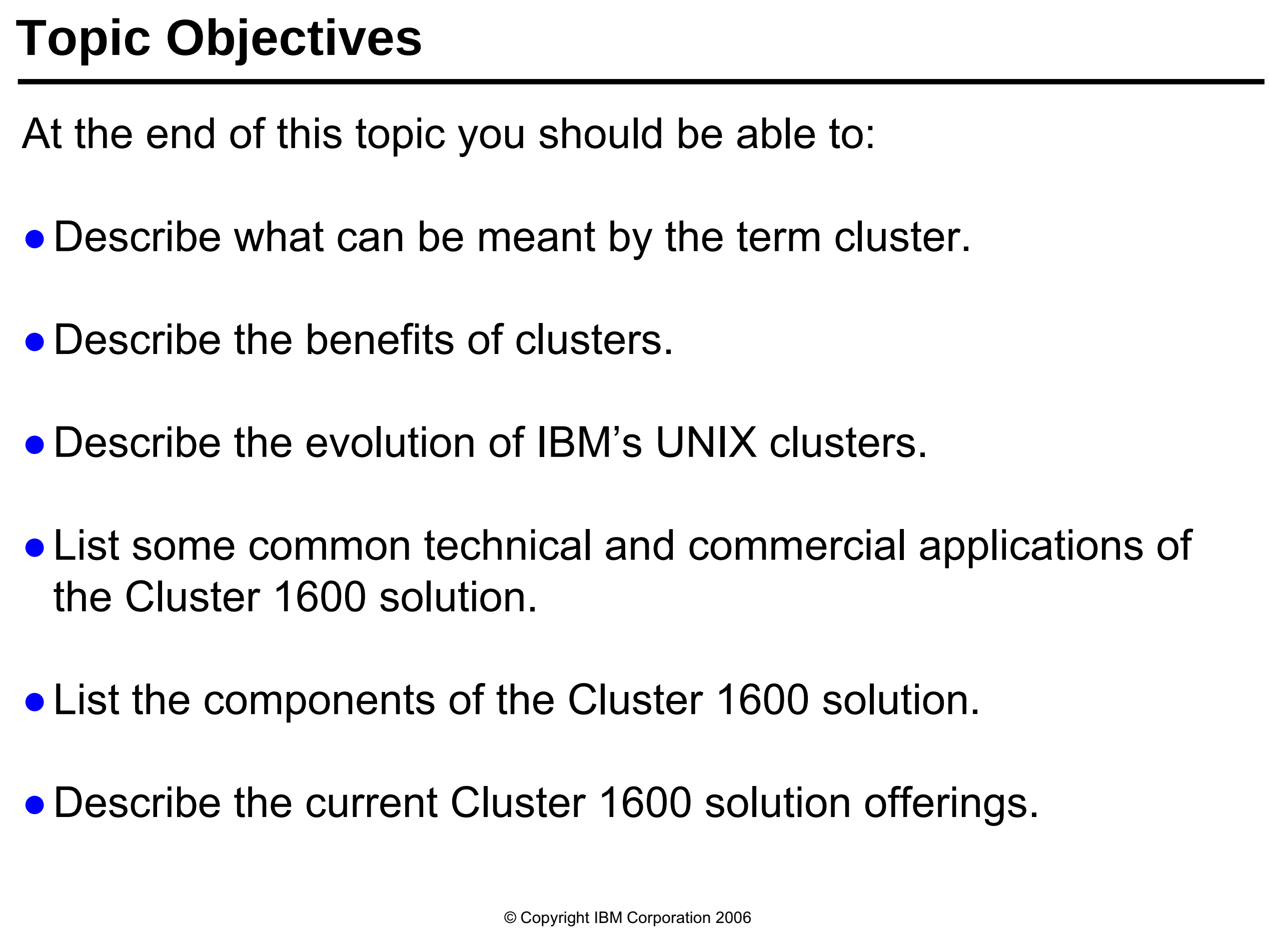

Topic Objectives

At the end of this topic you should be able to:

- Describe what can be meant by the term cluster.
- Describe the benefits of clusters.
- Describe the evolution of IBM’s UNIX clusters.
- List some common technical and commercial applications of the Cluster 1600 solution.
- List the components of the Cluster 1600 solution.
- Describe the current Cluster 1600 solution offerings.

Figure 5-27. Topic Objectives AM711.0

Notes:

More Buzzwords!

- Cluster
  - Group of compute nodes
- Node
  - A server
  - An LPAR
  - A blade
- High-speed interconnect
  - The cluster communications path

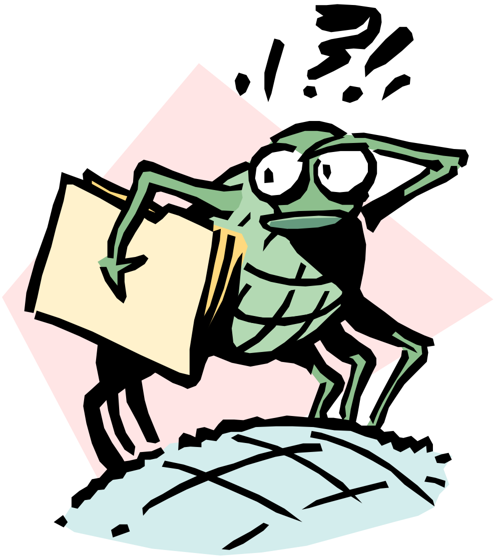

Figure 5-28. More Buzzwords! AM711.0

Notes:

Features of Cluster 1600

Scalability

Flexible architecture

System management

High availability

Figure 5-29. Features of Cluster 1600 AM711.0

Notes:

The Benefits of Clusters

Resource Consolidation

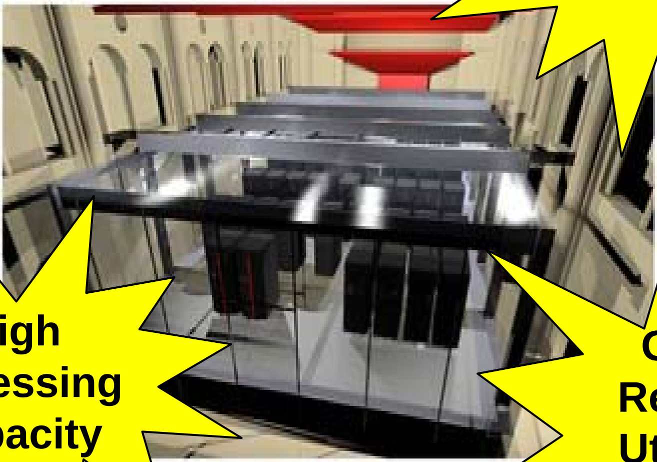

High Processing Capacity

Optimal Resource Utilization

Figure 5-30. The Benefits of Clusters AM711.0

Notes:

Clusters Can Provide:

- Geographical server consolidation
- 24 x 7 availability with failover protection
- Disaster recovery
- Scale-out and scale-up without downtime
- Centralized system management

Figure 5-31. Clusters Can Provide: AM711.0

Notes:

High Performance Computing (HPC) Clusters

- Engineering analysis
- Seismic
- Weather
- Computational chemistry and physics

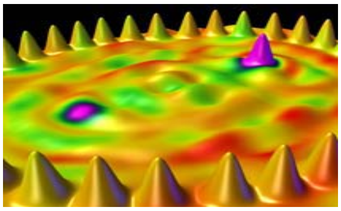

Figure 5-32. High Performance Computing (HPC) Clusters AM711.0

Notes:

Cluster for Simplified Management

- Large-capacity data and transaction volumes
- Scale-up (add processors) or scale-out (add servers) without downtime
- Single point-of-control for distributed and clustered server management
- Simplified use of IT resources
- 24x7 access to data and applications
- Business continuity in the event of disaster

Figure 5-33. Cluster for Simplified Management AM711.0

Notes:

Cluster 1600 Commercial Application Areas

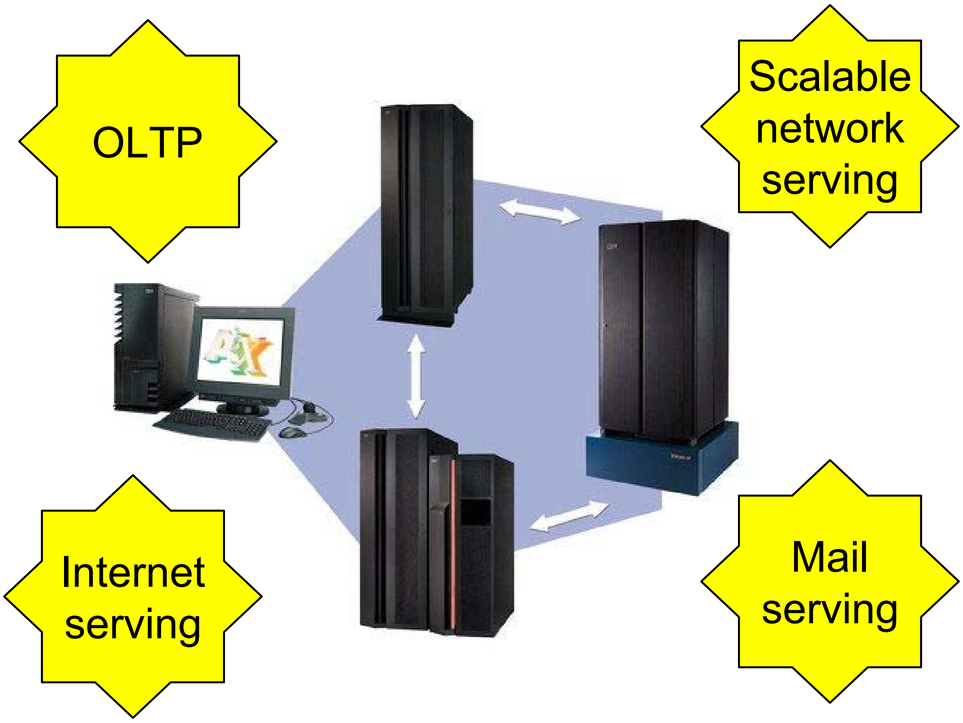

Figure 5-34. Cluster 1600 Commercial Application Areas AM711.0

Notes:

Cluster 1600 Components

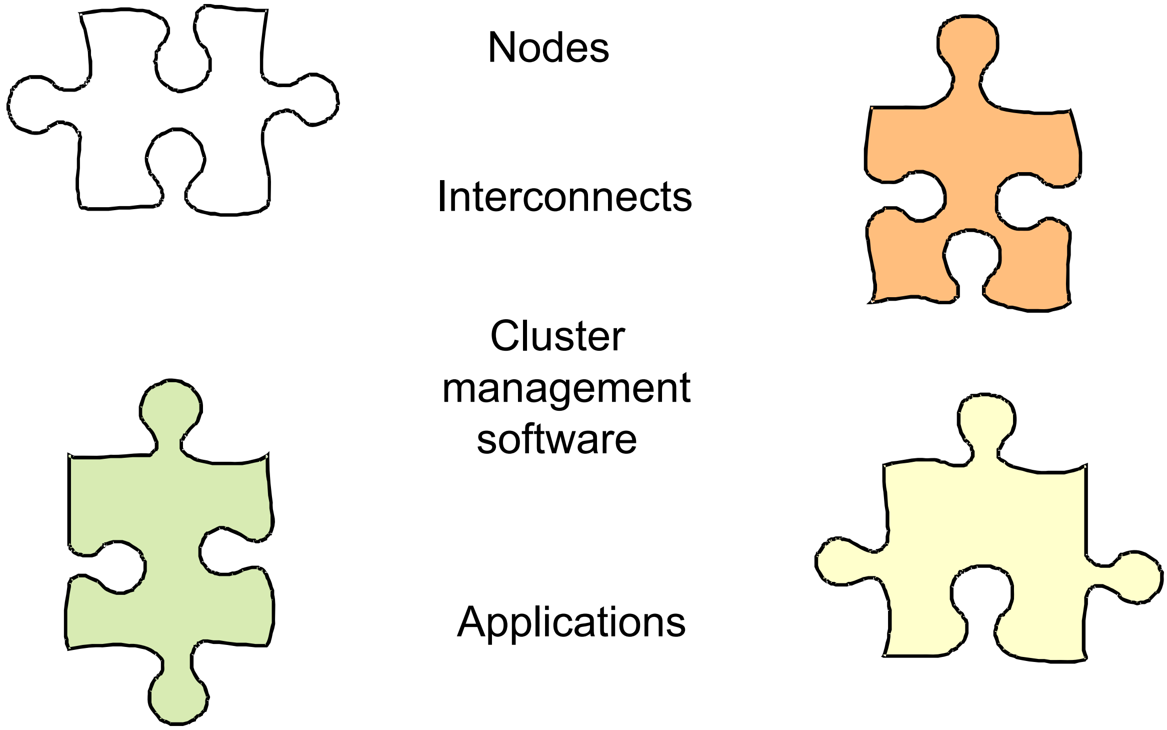

Nodes

Interconnects

Cluster
management
software

Applications

Figure 5-35. Cluster 1600 Components

AM711.0

Notes:

With the Cluster 1600 offerings, selected IBM POWER5™ and POWER4™ nodes can be logically partitioned (LPAR) by using Hardware Management Consoles (HMC), connected with industry-standard interconnects or IBM high-speed switches* and managed with IBM Cluster Systems Management (CSM) for AIX 5L™ and CSM for Linux® on POWER™ software. This combination of hardware and software provides powerful technology to help meet business needs.

Cluster 1600 Nodes

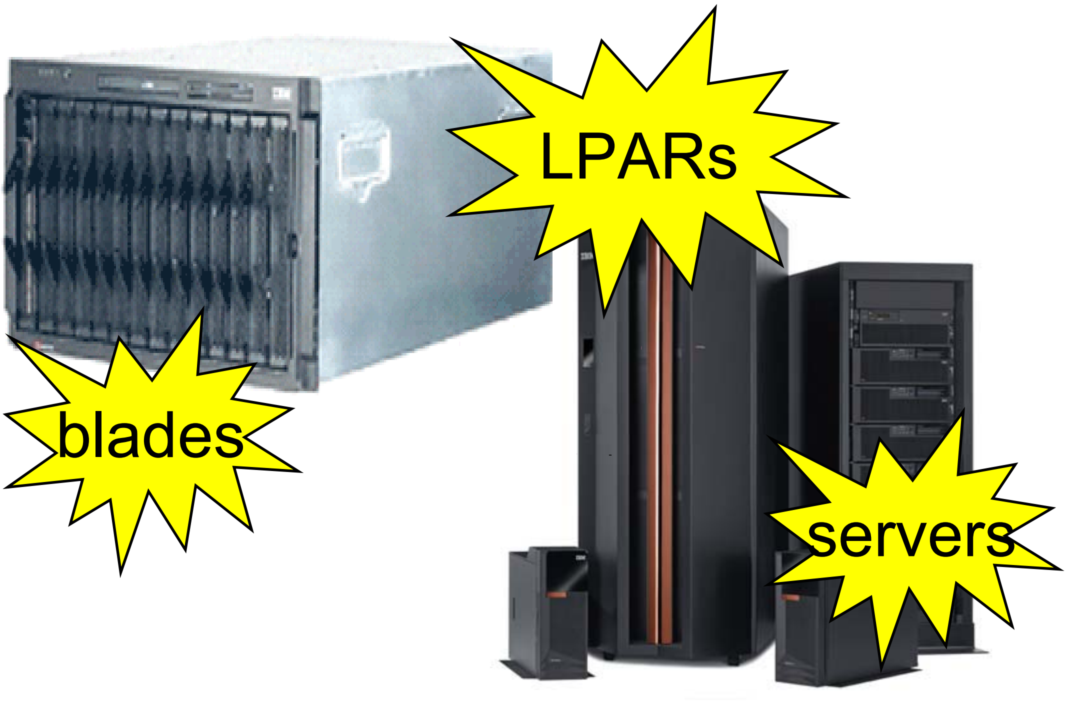

Figure 5-36. Cluster 1600 Nodes AM711.0

Notes:

POWER4 and POWER5 processor-based nodes cannot be intermixed on a switched network cluster.

Cluster 1600 Interconnects

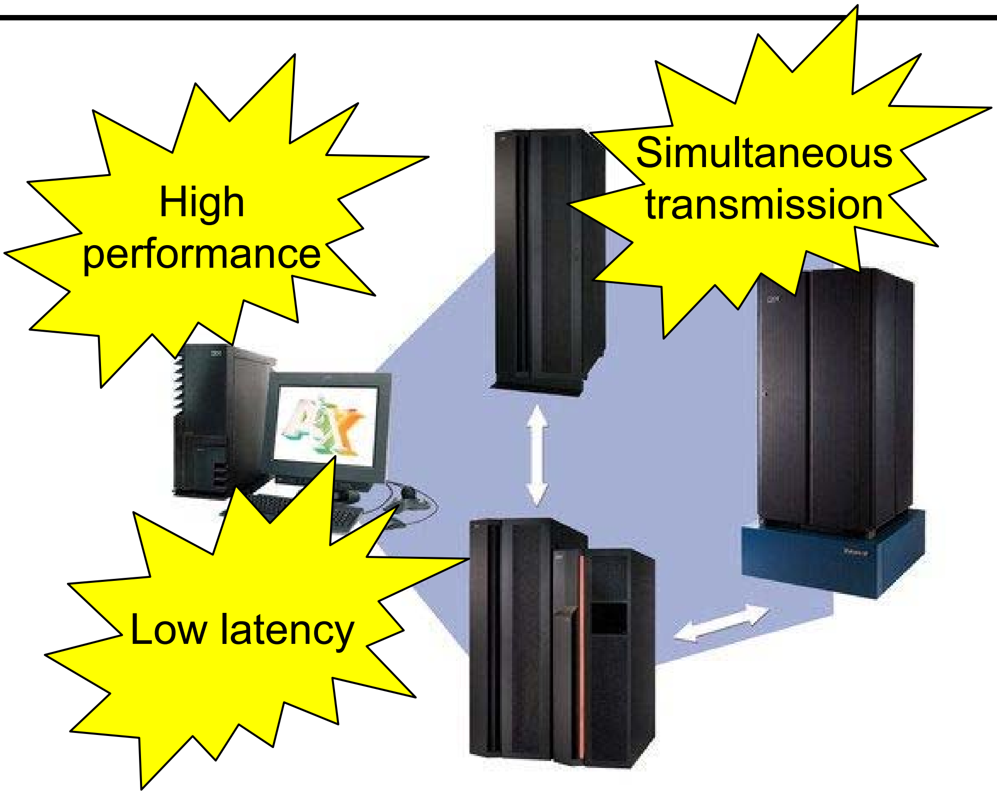

Figure 5-37. Cluster 1600 Interconnects AM711.0

Notes:

IBM High Performance Switch (HPS)

- 24 inch rack mount
- Supported by AIX 5L V5.2 and V5.3
- For p5-575, p5-590, p5-595
- For p655 and p690
- Mixing POWER4 and POWER5 nodes in the same HPS network is not supported

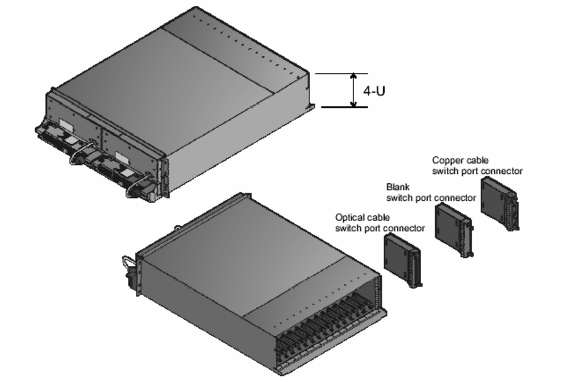
4-U
Copper cable switch port connector
Blank switch port connector
Optical cable switch port connector

Figure 5-38. IBM High Performance Switch (HPS) AM711.0

Notes:

CSM requires Ethernet connectivity for systems management and node installation.

The HPS is supported by AIX 5L V5.2 and V5.3.

The HPS topology imposes further limits on cluster sizes, detailed below.

General limits for p5-575, p5-590 and p5-595 clusters

Maximum switched operating system images 128 for p5-575; 16 for p5-590 and p5-595

Maximum servers per HMC 32

Maximum switch links in one-network configurations 128

Maximum switch links in two-network configurations 256

Note that mixed clusters including both POWER4 and POWER5 systems on the same HPS network are not supported.

Myrinet-2000 Switches

- Use when Linux is required operating system
- Popular in HPC clusters

JS20

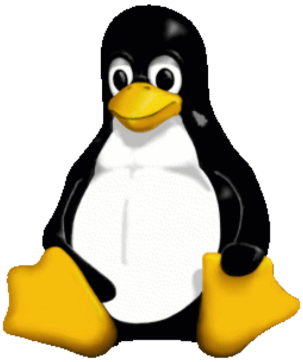

Myrinet Switches

Figure 5-39. Myrinet-2000 Switches AM711.0

Notes:

InfiniBand (IB) HCAs and Switches

- InfiniBand
  - Is an open standards-based high-bandwidth fabric
  - Permits high-speed low-latency interconnection between cluster servers
- IBM GX dual-port IB HCA (Galaxy)
  - 4x and 12x support (10 Gbs, 30 Gbs)
  - Higher performance
  - Virtualization across LPARs
  - Target P5 systems that feature external GX+ connector
- IBM PCI-X dual-port 4x IB HCA (Topspin)
  - OEM'd from Topspin/CISCO
  - Single solution across all system types
  - Target low-end P5 and Legacy P4 systems

Figure 5-40. InfiniBand (IB) HCAs and Switches AM711.0

Notes:

Industry standard Clustering Interconnect for Servers

Created by InfiniBand Trade Association (IBTA)

Organization formed by merger of Future I/O and Next Generation I/O

Members: Agilent, Dell, Hewlett-Packard, IBM, InfiniCon, Intel, Mellanox, Network Appliance, Oracle, Sun, Topspin, Voltaire

Characteristics,

Channel-based switched fabric

Point-to-point interconnect

Scalable with high performance

Cluster 1600 Standard Interconnects

- 10/100 Mbps Ethernet
- 2-port Gigabit Ethernet Fiber

Figure 5-41. Cluster 1600 Standard Interconnects AM711.0

Notes:

Cluster 1600 Software

- Cluster management software
  - CSM
  - PSSP
- Operating system
  - AIX
  - Linux
- Applications
  - HACMP
  - High performance computing software stack
  - GPFS

Figure 5-42. Cluster 1600 Software AM711.0

Notes:

Use Cluster System Management Software to:

- Install and update software on nodes
- Automate security configuration
- Run distributed commands
- Distribute files across cluster
- Monitor nodes
- Hardware control of nodes
- Manage node groups
- Run diagnostic tools

Figure 5-43. Use Cluster System Management Software to: AM711.0

Notes:

Cluster Systems Management Choices

- CSM when requirement is:
  - POWER5
  - AIX5L V5.3
  - Linux
- PSSP when requirement is existing client with:
  - SP nodes
  - SP switches

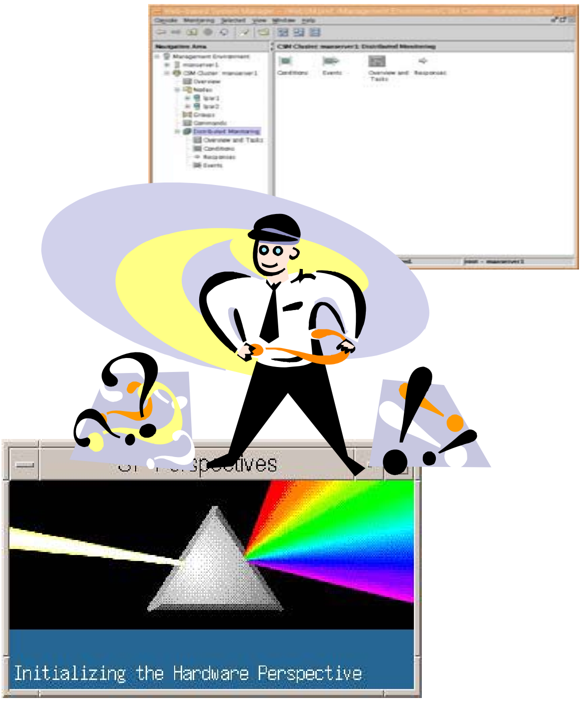

Figure 5-44. Cluster Systems Management Choices AM711.0

Notes:

AIX 5L V5.3 support

Supports POWER5 systems and IBM ~ BladeCenter™ JS20

New cluster functions

New option: High Availability Management Server (HAMS)

Copy files across nodes

Enablement of Kerberos 5

Usability enhancements

Additional Linux on POWER support

HA MS

System p5, ~ p5, pSeries and JS20

RHEL AS 3 and 4 support, SLES 9

High Performance Software

- LoadLeveler
- Parallel environment for AIX
- Engineering and Scientific Subroutine Library (ESSL) for AIX and Linux
- Parallel ESSL for AIX and Linux

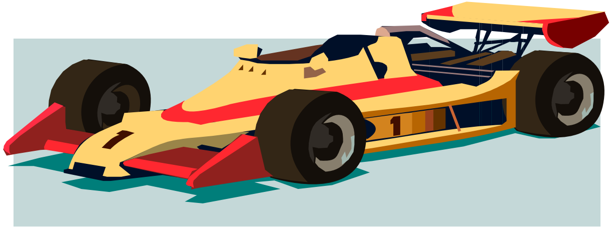

Figure 5-45. High Performance Software AM711.0

Notes:

GPFS – General Parallel File System

- Among the highest performing file systems in the world
- Powers several of the world's most powerful supercomputers
- Extremely reliable – eliminates all single points of failure
- Applicable to many commercial applications:
  - Digital media,
  - Business intelligence
  - Financial analysis

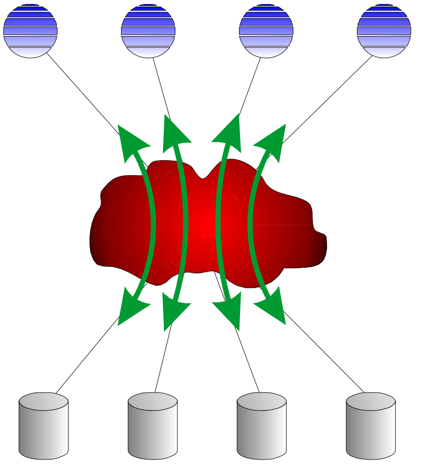

Figure 5-46. GPFS – General Parallel File System AM711.0

Notes:

General Parallel File System is an extremely reliable, high-performance file system that powers many of the world's most powerful research supercomputers. GPFS is now finding its way into commercial customers for two reasons. The first is that as data volumes and processor speeds increase, disk I/O has become a bottleneck; GPFS is an excellent solution if standard network or SAN file systems cannot deliver the aggregate bandwidth that your customers' applications require. The second is its reliability – even some customers that don't need the performance have implemented GPFS because it can eliminate the bottlenecks and single points of failure inherent in standard file systems. GPFS is finding a growing audience in digital media, financial services, business intelligence, life sciences, petrochemicals and other data-intensive applications.

High Availability Clusters

- When application availability is important
- Peer cluster topology
- Nodes communicate and monitor each other
- High availability cluster multiprocessing software (HACMP)

Figure 5-47. High Availability Clusters AM711.0

Notes:

If node A fails, HACMP can detect the failure and migrate the web server to node B to maintain the application availability. When the problem is resolved, HACMP can migrate the web server back to node A or leave the application running as-is.

The concern with this type of failover is performance because the work of two machines is being handled by one; HACMP clusters can scale up to 32 nodes to accommodate high-traffic business applications. With appropriate servers, active-active configurations can provide business resilience to customers who refuse to purchase redundant, predominantly idle hardware.

HACMP 5.3

- Additional ease of configuration and management enhancements
- Flexible resource group configuration
- Online planning worksheets in XML format
- Support for more OEM volume groups and filesystems including:
  - Veritas volume manager and Veritas file system
- Support for InfiniBand adapters

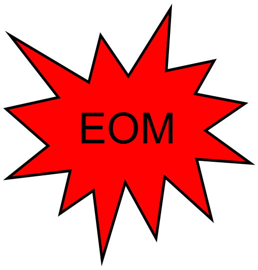

Figure 5-48. HACMP 5.3 AM711.0

Notes:

HACMP 5.3 Even Simpler

Online Planning Worksheets **(OLPW)** cluster definition saved in well-defined **XML format** for easy edits. Cluster snapshot can also be saved to this format.

Accommodation of **OEM volume groups and file systems** (in addition includes out-of-the-box support for Veritas Volume Manager (VxVM) and VxFS using **Veritas** Foundation Suite v. 4.0).

SMS pager capability for remote text or cell phone notifications

HACMP V5.4

New July 2006

- Web-based GUI
- Improved cluster verification tools
- Smart Assists for DB2, Oracle, and WebSphere
- Non-disruptive HACMP cluster startup, upgrades, and real-time maintenance
- Fast Failure Detection and takeover of production applications on backup servers
- Metro Mirror support for intermixed environments (DS6000, DS8000, ESS800)
- HACMP/XD GLVM Multi-Link feature for improved data mirroring protection and performance
- Concurrent mode access for simultaneous applications execution at local site
- The optional features HACMP/XD and HACMP Smart Assist provide *complete* high availability disaster recovery solutions for your business on AIX 5L.

Figure 5-49. HACMP V5.4 AM711.0

Notes:

HACMP for Linux V5.4

- Same user interface and configuration as HACMP for AIX 5L
- Across-the-board support for System p and System i running Linux
- Support for clusters of up to eight nodes
- Application monitoring and failover
- Network/IP interface monitoring and failover
- Node monitoring and failover
- Web-based GUI configuration and status display
- Most of the function of basic HACMP for AIX 5L, with the exception of shared storage
- Proven HACMP technology

New July 2006

Figure 5-50. HACMP for Linux V5.4 AM711.0

Notes:

HACMP Extended Distance (XD)

- Disaster recovery (remote failover) is becoming a major customer focus
- The HACMP/XD option provides a basket of disaster recovery capabilities and integration points
  - Includes both IP-based mirroring and integration for hardware-based mirroring
- Backup systems may be physically separate from primary operations for protection in the event of power failure, flood, earthquake, and so forth

Figure 5-51. HACMP Extended Distance (XD) AM711.0

Notes:

Cluster 1600 Scalability Limits

Figure 5-52. Cluster 1600 Scalability Limits AM711.0

Notes:

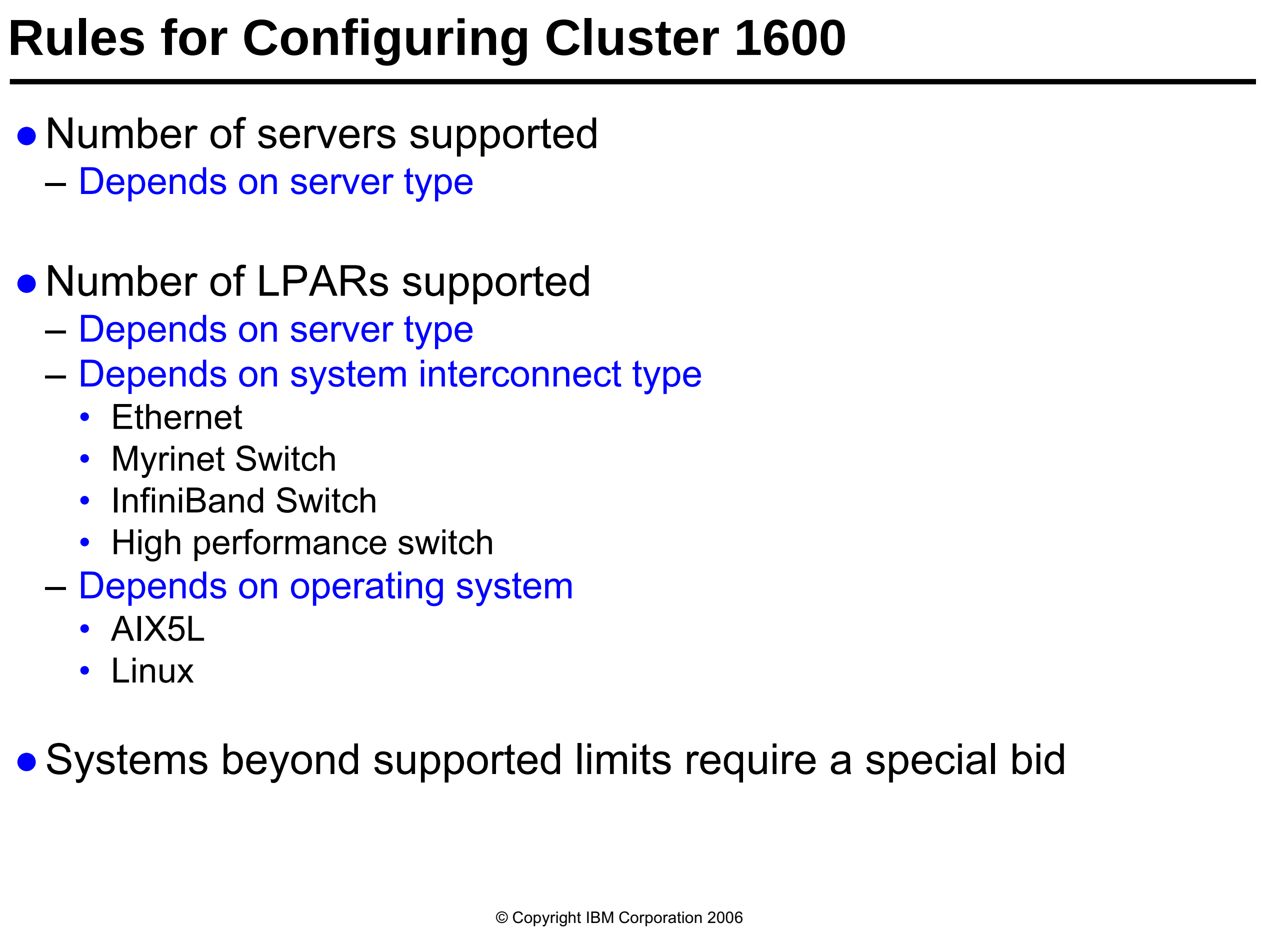

Rules for Configuring Cluster 1600

- Number of servers supported
  - Depends on server type
- Number of LPARs supported
  - Depends on server type
  - Depends on system interconnect type
    - Ethernet
    - Myrinet Switch
    - InfiniBand Switch
    - High performance switch
  - Depends on operating system
    - AIX5L
    - Linux
- Systems beyond supported limits require a special bid

Figure 5-53. Rules for Configuring Cluster 1600 AM711.0

Notes:

Cluster 1600 Configuration Rules

- General AIX5L with CSM rules

1.	Maximum LPARs per cluster	512
2.	Maximum servers per HMC	32
3.	Maximum LPARs per HMC	128

Figure 5-54. Cluster 1600 Configuration Rules AM711.0

Notes:

A Cluster 1600 with CSM for AIX 5L can consist of between 2 and 512 AIX5L operating system images; specific limits are determined by the type of server and interconnect (Ethernet, InfiniBand or IBM HPS) being used. Any cluster system exceeding any of the announced support limits requires a special bid. An operating system image (OSI) or logical node can be a full-system partition or a logical partition (LPAR) on those servers that support LPARs.

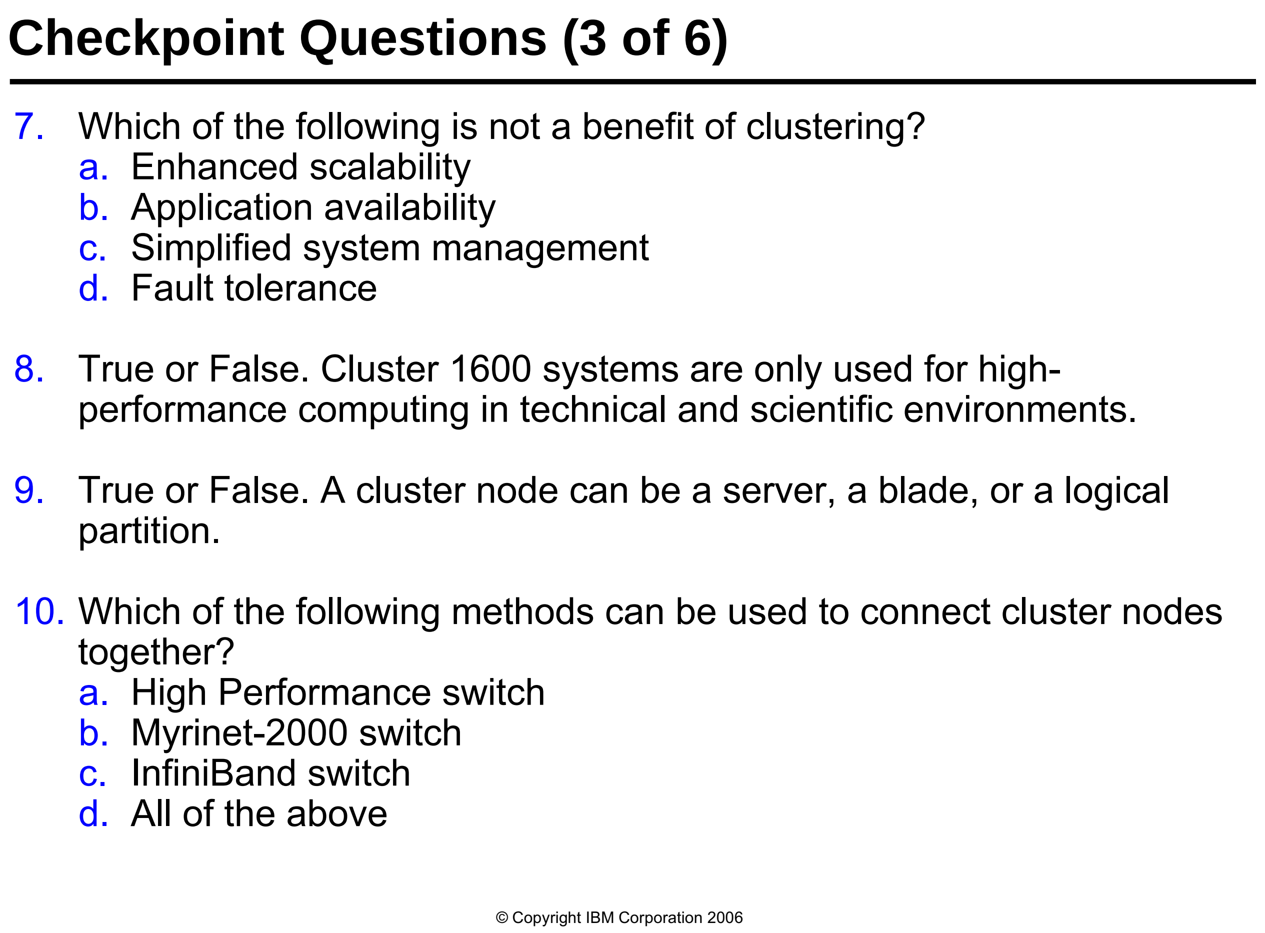

Checkpoint Questions (3 of 6)

7. Which of the following is not a benefit of clustering?
    a. Enhanced scalability
    b. Application availability
    c. Simplified system management
    d. Fault tolerance

8. True or False. Cluster 1600 systems are only used for high-performance computing in technical and scientific environments.

9. True or False. A cluster node can be a server, a blade, or a logical partition.

10. Which of the following methods can be used to connect cluster nodes together?
    a. High Performance switch
    b. Myrinet-2000 switch
    c. InfiniBand switch
    d. All of the above

Figure 5-55. Checkpoint Questions (3 of 6) AM711.0

Notes:

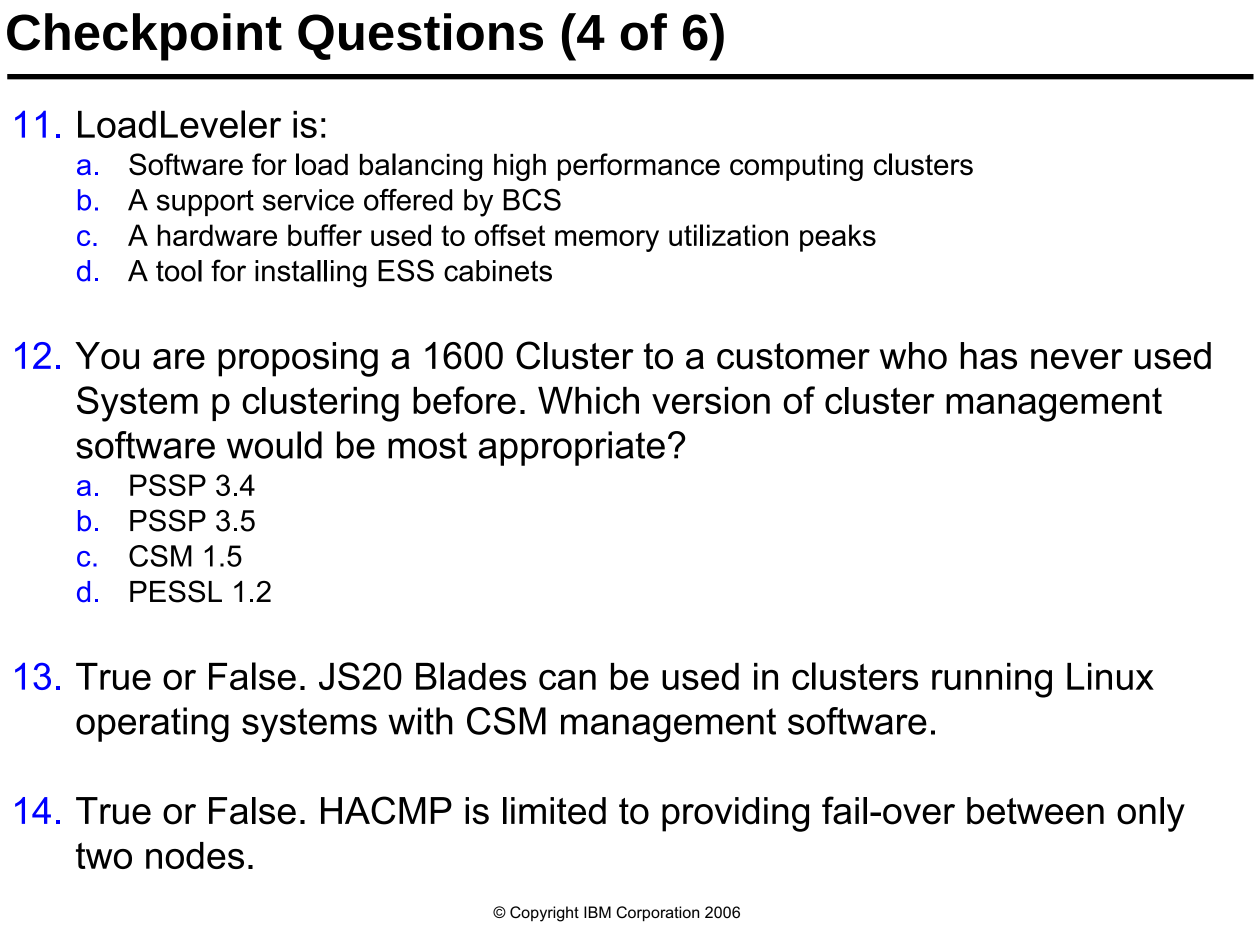

Checkpoint Questions (4 of 6)

11. LoadLeveler is:
    - a. Software for load balancing high performance computing clusters
    - b. A support service offered by BCS
    - c. A hardware buffer used to offset memory utilization peaks
    - d. A tool for installing ESS cabinets

12. You are proposing a 1600 Cluster to a customer who has never used System p clustering before. Which version of cluster management software would be most appropriate?
    - a. PSSP 3.4
    - b. PSSP 3.5
    - c. CSM 1.5
    - d. PESSL 1.2

13. True or False. JS20 Blades can be used in clusters running Linux operating systems with CSM management software.

14. True or False. HACMP is limited to providing fail-over between only two nodes.

Figure 5-56. Checkpoint Questions (4 of 6) AM711.0

Notes:

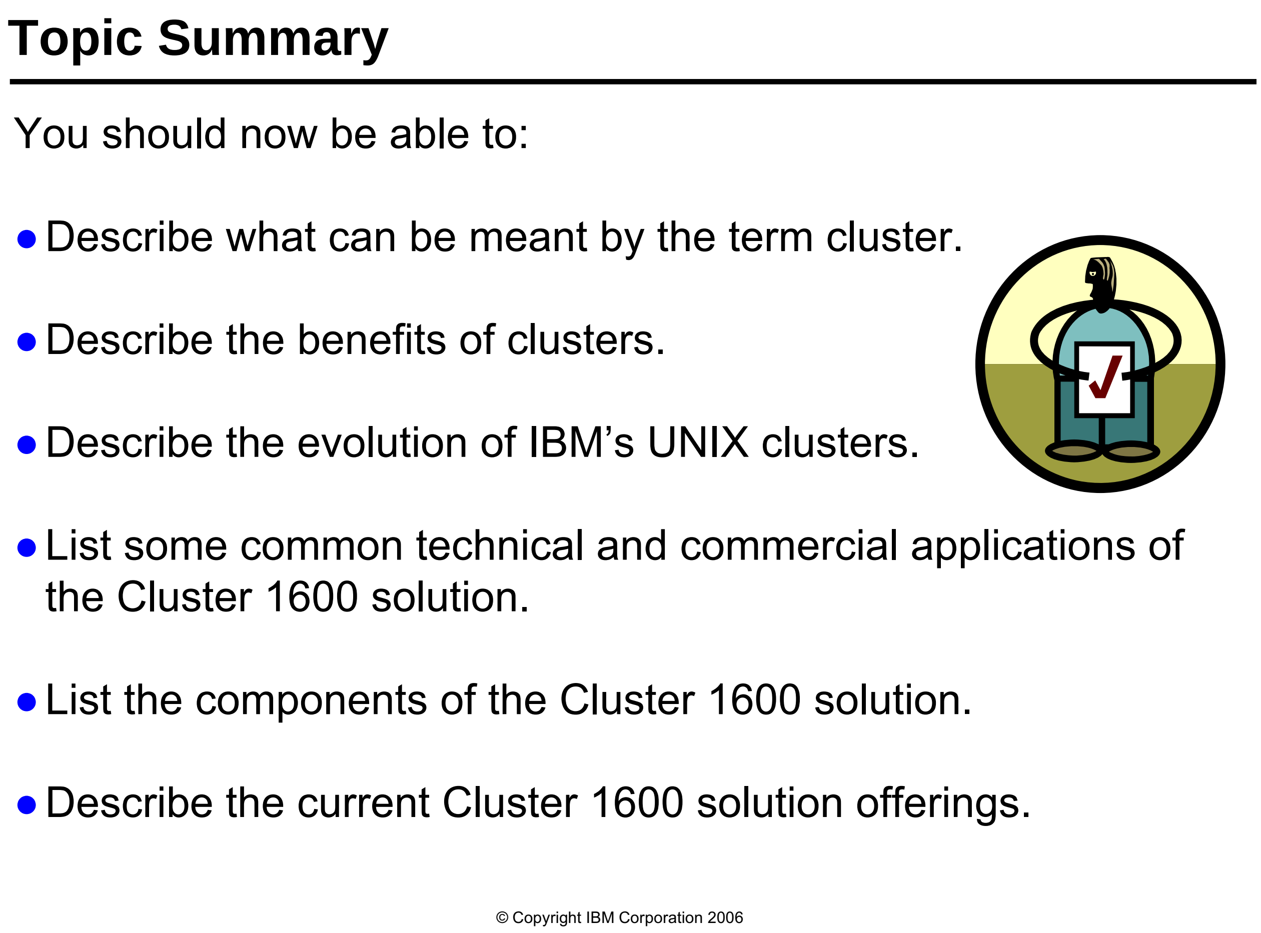

Topic Summary

You should now be able to:

- Describe what can be meant by the term cluster.
- Describe the benefits of clusters.
- Describe the evolution of IBM's UNIX clusters.
- List some common technical and commercial applications of the Cluster 1600 solution.
- List the components of the Cluster 1600 solution.
- Describe the current Cluster 1600 solution offerings.

Figure 5-57. Topic Summary AM711.0

Notes:

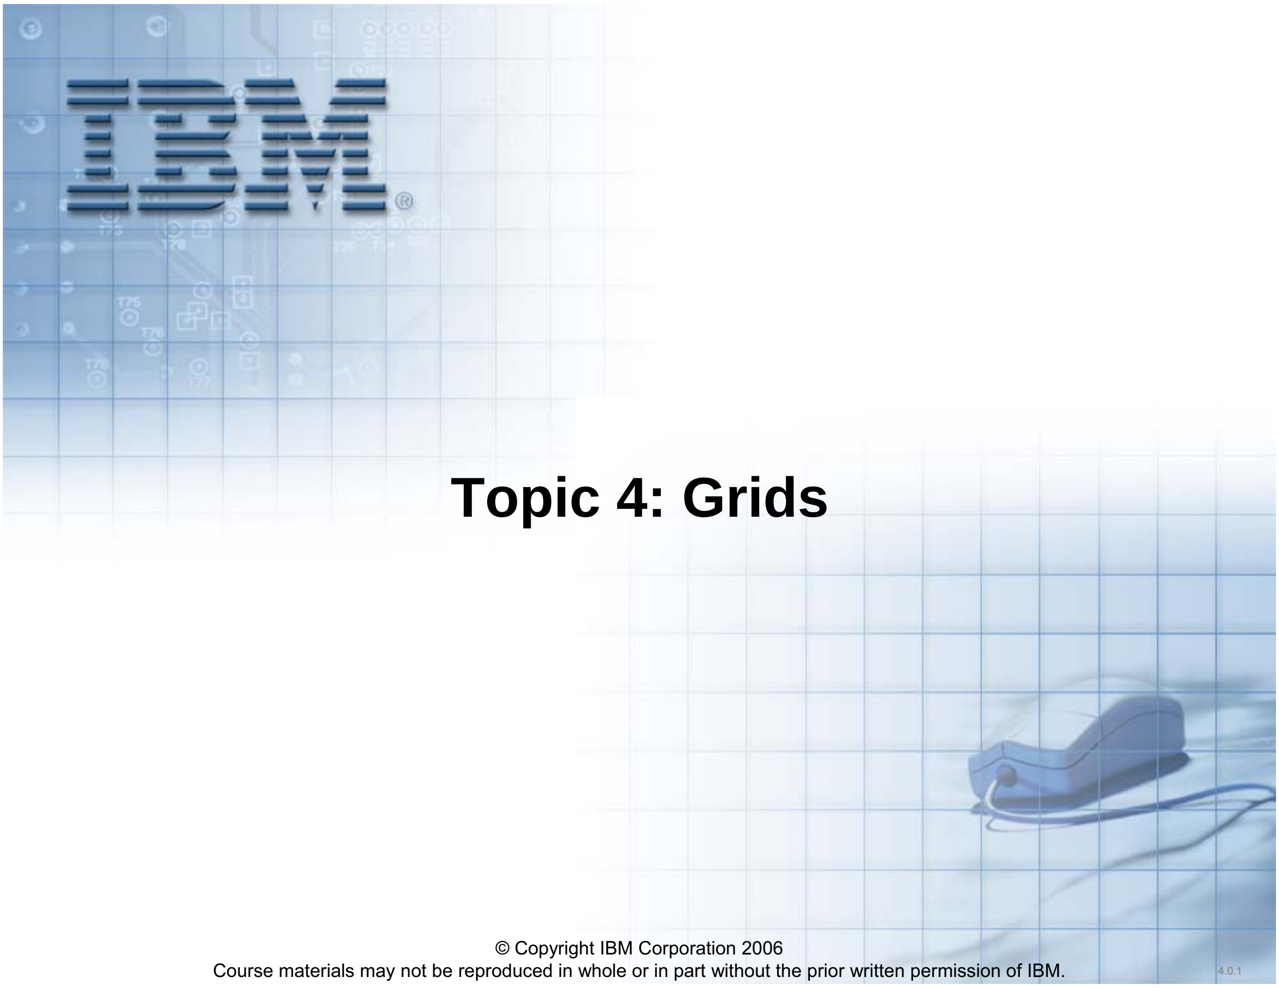

Topic 4: Grids

4.0.1

Figure 5-58. Topic 4: Grids AM711.0

Notes:

Topic Objectives

At the end of this topic you should be able to:

- Describe what Grid computing is.
- Know why and when to sell Grid computing.

Figure 5-59. Topic Objectives AM711.0

Notes:

So, What is Grid Computing?

- Distributed computing based on open standards for enabling heterogeneous resources
- A computing model that optimizes the productivity of your IT resources by maximizing utilization levels without a considerable increase in operating costs
  - It is all about business value
- A critical component of IBM’s On Demand vision

Figure 5-60. So, What is Grid Computing? AM711.0

Notes:

Grid Focus Areas and Industries

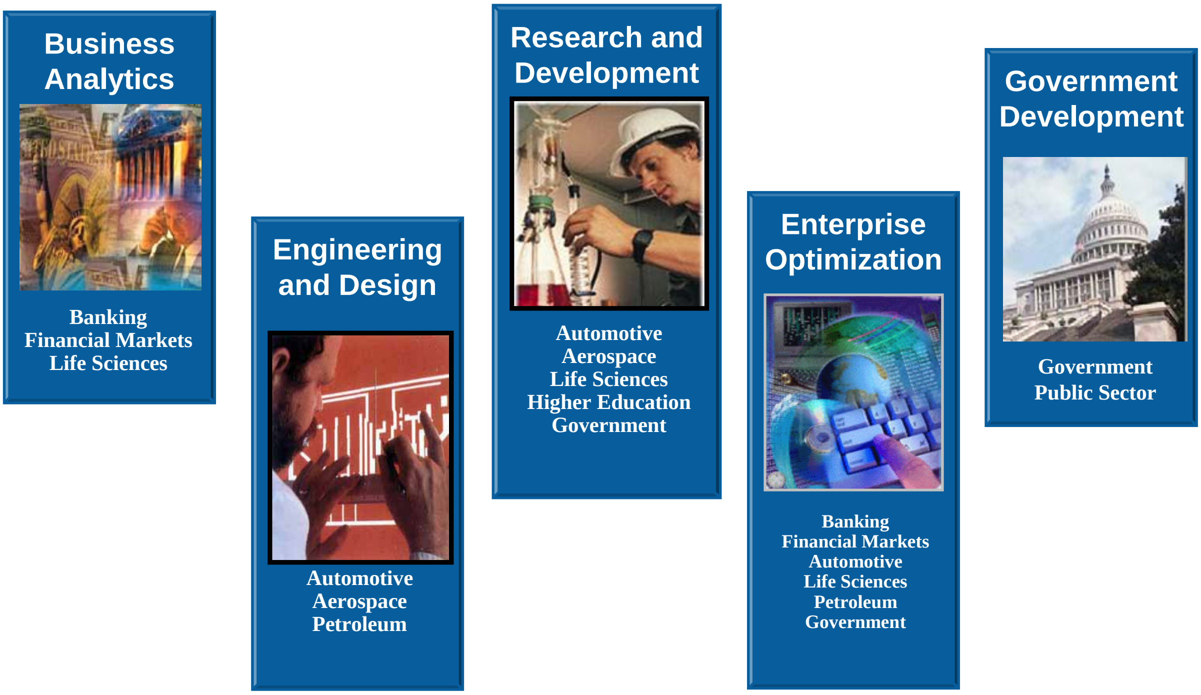

Business Analytics

Banking
Financial Markets
Life Sciences

Engineering and Design

Automotive
Aerospace
Petroleum

Research and Development

Automotive
Aerospace
Life Sciences
Higher Education
Government

Enterprise Optimization

Banking
Financial Markets
Automotive
Life Sciences
Petroleum
Government

Government Development

Government
Public Sector

Figure 5-61. Grid Focus Areas and Industries

AM711.0

Notes:

Why Should You Sell Grid Computing?

- Helps you sell more stuff by…
  - Improving server and storage utilization by virtualizing those resources and then applying those resources against common problems.
  - Reducing application runtimes through Grid enablement (for example, job and task scheduling across distributed, heterogeneous resources).
  - Creating high availability and improve uptime.
- Opens doors to new clients
- Quickly creates business and IT value for your clients
- Grid is always on demand

Figure 5-62. Why Should You Sell Grid Computing? AM711.0

Notes:

Grid and its value proposition is a game changing venue. To customers it is something that is “new” and fresh. This enables the TotalStorage sales team an opportunity to show customers how their portfolio delivers that broadest possible range of solutions that address the needs that Grid places on the “information infrastructure.” The key is for our sales teams to know that they have the most comprehensive set of solutions to satisfy the needs of building an “information infrastructure” for Grid computing.

How Grid Drives Server Solutions

- Your customers are interested in Grid technologies
- Grid computing increases the need for new HW over time
- Encourages more complex IT operations
- Creates a use for excess/unused capacity
- Creates a need for Grid-enabled HW and SW
- Many Grid deployments accompany a shift from UNIX to Linux

Figure 5-63. How Grid Drives Server Solutions AM711.0

Notes:

First, customers bought hardware from Sun during the dotcom Internet build out because Sun was the leader, they were setting the agenda. Now EMC is starting to play the Grid card. This time we are setting the agenda and our leadership in Grid will set the stage for many opportunities, not just a single hardware sale.

Grid computing requires that a flexible “information infrastructure” be in place to enable applications to access data regardless of where the application gets deployed into the Grid. Moreover, beyond the needs of computing, collaboration within business organizations (and beyond the enterprise) are seeking solutions to enable data sharing.

Where Can I Learn More About Grid?

- Grid Web site - www.ibm.com/grid
  - Articles, press releases, analyst reports
  - Web casts, brochures, solutions briefs
  - Customer success stories
- Grid w3 site - w3.grid.ibm.com
  - SGB Seller Playbook
  - Education / Classes, FAQ, Pocket Guide
  - Demos, videos, key Grid contacts
- Redbook on Data and Information Grids w3.itso.ibm.com
  - Introduction to an Information Infrastructure for Grid Computing
  - Internal Use Only - IBM Publication Number: ZZ81-0486
- *NEW* SGB Seller Playbook - w3.grid.ibm.com
  - Includes Grid, VS, Deep and Linux
  - New and improved industry messages and customer materials
  - Seller resource guides, mapped to SSM steps
  - Qualification questions and next steps to take with your client

Figure 5-64. Where Can I Learn More About Grid? AM711.0

Notes:

Checkpoint Questions (5 of 6)

15. What is Grid computing?
    a. Distributed computing based on open standards for enabling heterogeneous resources.
    b. A computing model that optimizes the productivity of your IT resources by maximizing utilization levels without a considerable increase in operating cost.
    c. A critical component of IBM’s On Demand vision.
    d. All of the above.

16. One obstacle to enterprise-wide Grid computing is:
    a. Political issues between different lines of business.
    b. The technology to support this is not available.
    c. There are not sufficient accounting tools available.
    d. A lack of scalability of the Grid computing model.

Figure 5-65. Checkpoint Questions (5 of 6) AM711.0

Notes:

Checkpoint Questions (6 of 6)

17. Which of the following are benefits of Grid computing?
    a. Application Acceleration
    b. Aggregating Information
    c. Infrastructure Flexibility
    d. All of the above

18. Which of the following is a Grid computing focus area?
    a. Business Analytics
    b. Engineering and Design
    c. Enterprise Optimization
    d. All of the above

Figure 5-66. Checkpoint Questions (6 of 6) AM711.0

Notes:

Topic Summary

You should now be able to:

- Describe what Grid computing is.
- Know why and when to sell Grid computing.

Figure 5-67. Topic Summary AM711.0

Notes:

Unit Summary

Having completed this unit you should be able to:

- Discuss System p products which support our clients' requirements for infrastructure.
- List the benefits of networking.
- List the main features of System p network support.
- List the network adapters supported by System p.
- Identify appropriate adapters from customer requirements.
- Describe what can be meant by the term cluster.
- Describe the benefits of clusters.
- List some common technical and commercial applications of the Cluster 1600 solution.
- List the components of the Cluster 1600 solution.
- Describe the current Cluster 1600 solution offerings.
- Describe what Grid computing is.
- Know why and when to sell Grid computing.

Figure 5-68. Unit Summary AM711.0

Notes:

Unit 6. Storage

What This Unit Is About

This unit describes the IBM Storage products relevant to System p solution design.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Position IBM System Storage disk products
- Position IBM System Storage tape products
- Discuss IBM System Storage software solutions

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

After completing this unit, you should be able to:

- Position IBM System Storage disk products.
- Position IBM System Storage tape products.
- Discuss IBM System Storage software solutions.

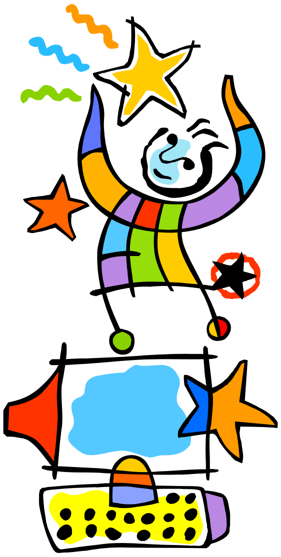

Figure 6-1. Unit Objectives

AM711.0

Notes:

The Storage Opportunity

- External disk market grew at 10% in 2005 against a projection of 4%
- IBM System Storage is gaining share against the competition
- The challenges facing our clients play to IBM's strengths

Figure 6-2. The Storage Opportunity AM711.0

Notes:

Storage Challenge

Identify the "Right" Data …

Put it on the "Right" Media …

At the "Right" Time …

… Until Ultimate Disposal / Destruction.

Figure 6-3. Storage Challenge AM711.0

Notes:

Unique Synergy IBM System Storage and System p

- Collaborate on architecture and design.
  - DS8000, DS6000, ESS and NAS 500 are built on Power Architecture technology
- System Storage is the muscle behind System p
  - TPC-C 3,210,540 tpmC* - DS4500 and p5-595
- Conduct early, rigorous and comprehensive stress testing in System p and System Storage labs.
- Share support experience and knowledge between System p and System Storage technical specialists.

* Current as of 10/04/05.   TPC-C result of 3,210,540 tpmC, $5.19/tpmC, configuration available on May 14, 2005

Figure 6-4. Unique Synergy IBM System Storage and System p AM711.0

Notes:

This helps IBM System Storage and pSeries development to:

- Better design products that work well together with more robust interlocking.
- Bypass potential pitfalls with fixes implemented before clients are ever impacted.
- Implement streamlined, efficient, integrated product offerings.
- Ensure that products that have passed stringent testing together.

The Benefits of this Synergy

- Helps IBM System Storage and System p developers to:
  - Design better products that work well together with more robust interlocking.
  - Bypass potential pitfalls with fixes implemented before clients are ever impacted.
  - Implement streamlined, efficient, integrated product offerings.
  - Ensure that products that have passed stringent testing together.

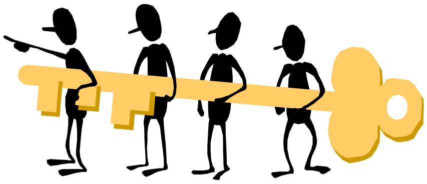

Figure 6-5. The Benefits of this Synergy AM711.0

Notes:

This provides value to pSeries and System Storage clients by helping to:

- Contribute to IBM Enterprise Disk continuous quality improvement.
- Verify product reliability even in heavy stress environments.
- Speed server and storage implementation.
- Improve infrastructure investments.
- Avoid problems.

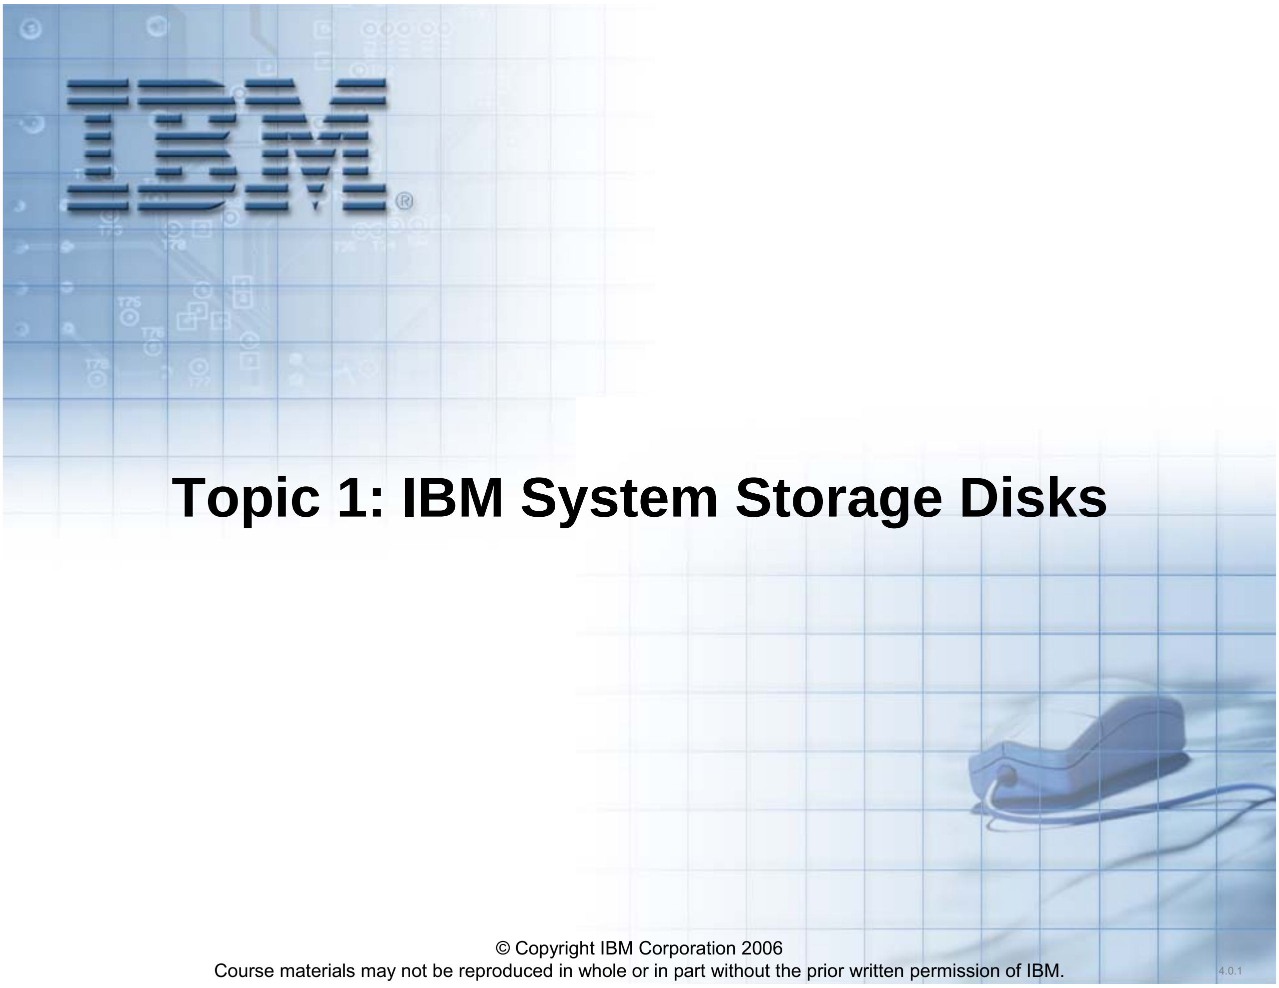

Topic 1: IBM System Storage Disks

4.0.1

Figure 6-6. Topic 1: IBM System Storage Disks AM711.0

Notes:

Benefits of IBM TotalStorage Disk Systems

IBM TotalStorage disk system innovations are designed to help our clients:

- **Simplify** the underlying IT infrastructure of storage and its management to lower cost and complexity while increasing the ability to respond to changing needs.
- Maintain **business continuity**, security and data durability.
- Efficiently manage information **throughout its lifecycle**, relative to its business value.

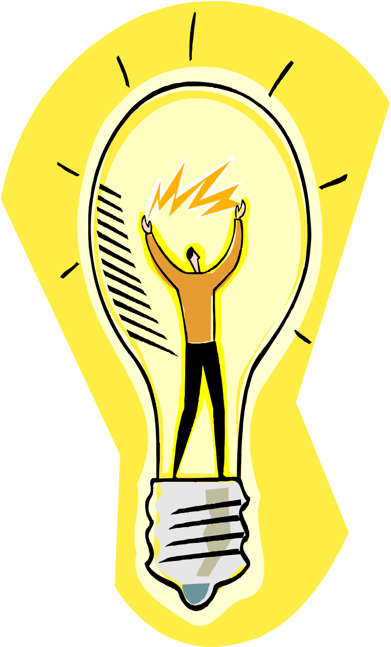

Figure 6-7. Benefits of IBM TotalStorage Disk Systems AM711.0

Notes:

Storage Buzzwords, Acronyms, and TLAs

- What is Fibre Channel?
- What is a SAN?
- What is a SAN fabric?
- What is NAS?
- What is iSCSI?
- What is RAID?
- What is TOE?

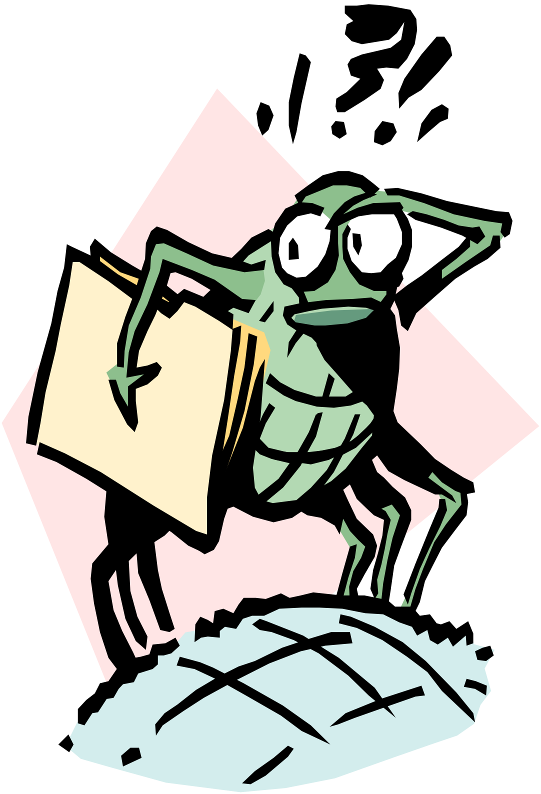

Figure 6-8. Storage Buzzwords, Acronyms, and TLAs AM711.0

Notes:

IBM System Storage Disk Continuum

- Broad range of offerings
- Entry-level to Enterprise
- DASD, SAN, and NAS

Figure 6-9. IBM System Storage Disk Continuum

AM711.0

Notes:

IBM System Storage EXP24 (7031-D24, 7031-T24)

- 7031-D24
  - 4U 19” Rack model
- 7031-T24
  - Tower model
- Up to 24 Ultra320 hot-swappable disk drives
- New 15 K rpm 146.8 GB disks

Rack model

Tower model

Figure 6-10. IBM System Storage EXP24 (7031-D24, 7031-T24) AM711.0

Notes:

Disk options:

73.4 / 146.8 / 300GB	10K rpm
36.4 / 73.4 / 146.8GB	15K rpm

Four independent disk drive 6-packs

- Single or Dual Bus repeater cards for each 6-pack
- Dual Bus repeater cards for high availability
- Optionally configure as two 12-packs

Redundant cooling and optional redundant power

PCI-X DDR Dual Channel Ultra 320 SCSI Adapters

- Support for data transfer speed of 320 MBs
- Dual, independent LVD buses.
- Support for a variety of storage offerings including:
  - System Storage EXP24 Ultra320 SCSI Expandable Disk enclosure
- RAID adapter provides RAID 0,5,6 or 10 capability
- Address up to 30 16-bit SCSI physical disk drives on two separate buses

Daughter card

PCI-X adapter

Figure 6-11. PCI-X DDR Dual Channel Ultra 320 SCSI Adapters AM711.0

Notes:

Entry Level Storage Solutions

IBM TotalStorage DS400

- 14-Drive, 2 Gb FC to U320 SCSI subsystem
- Single and dual RAID controller options
- Scalable to 5.8 TB (40 drives with 2 EXP400)
- Common software, RAID code, and drive carrier with DS300

IBM TotalStorage DS300

- 14-Drive, 1 Gbe iSCSI to U320 SCSI subsystem
- Single and dual RAID controller options
- Scalable to 2 TB (14 drives)
- Common software, RAID code and drive carrier with DS400

Figure 6-12. Entry Level Storage Solutions AM711.0

Notes:

PCI-X 1Gbps iSCSI TOE Adapters

- TCP/IP Offload Engine (TOE)
  - Reduce system processor utilization by encapsulating SCSI commands and data into TCP packets.
- Copper or fiber optical interface
- Connection direct to storage or to SAN or NAS storage systems such as N3000 and N5000 iSCSI enabled offerings
- Lower cost alternative than Fibre Channel for connection to SANs

Figure 6-13. PCI-X 1Gbps iSCSI TOE Adapters

AM711.0

Notes:

IBM System Storage N3000 – IP-Attached Storage

- Customers who need Windows, Linux, and UNIX file consolidation, and storage
- Consolidate storage from older storage devices
- Support mixed user environments
- Single Storage Infrastructure for NAS and iSCSI environments
- Initial focus is on SMB and distributed environments

- N3700 available in single or dual storage controller
- Compact 3U drawer scales to 4 TB
- Enabler for IP-based SANs
- Rock solid foundation offers redundant hot swappable components and robust copy services

Figure 6-14. IBM System Storage N3000 – IP-Attached Storage AM711.0

Notes:

N3700 available in single or dual storage controller

Cluster controller offers redundancy, higher availability and increased performance

Compact 3U drawer scales to 4TB

16TB maximum configuration with expansion units

Single platform offers dual dialect support

NAS and iSCSI protocol attachment to Windows, UNIX and Linux environments

Suitable for lights-out operation in distributed environments

Extremely suitable for SMB mission-critical environments

IBM System Storage N5000 – IP-Attached Storage

- N5000 series (N5200 and N5500)
  - Announced: November 8, 2005
  - General Availability: December 9, 2005
- Newest offerings utilizing Network Appliance technology
  - Equivalent to FAS3020 and FAS3050
- Concurrently supports NAS, IP SAN, and FC SAN environments
- Serves three highly desired environments
  - High-performance, mission-critical data environments
  - Archive and data retention environments
  - Most important, disaster recovery site missions

Figure 6-15. IBM System Storage N5000 – IP-Attached Storage AM711.0

Notes:

IBM System Storage DR550 and DR550 Express

- Preconfigured, integrated hardware and software
- Data protection options such as encryption and user authentication
- Affordable, high-performance offering
- Comprehensive suite of software features for policy-based and event-based data management
- Data encryption for disk and all optionally attached storage devices (for example, tape)
- Automatic provisioning, migration, expiration, and archiving capabilities
- Ability to use advanced tape to back up data objects

Figure 6-16. IBM System Storage DR550 and DR550 Express AM711.0

Notes:

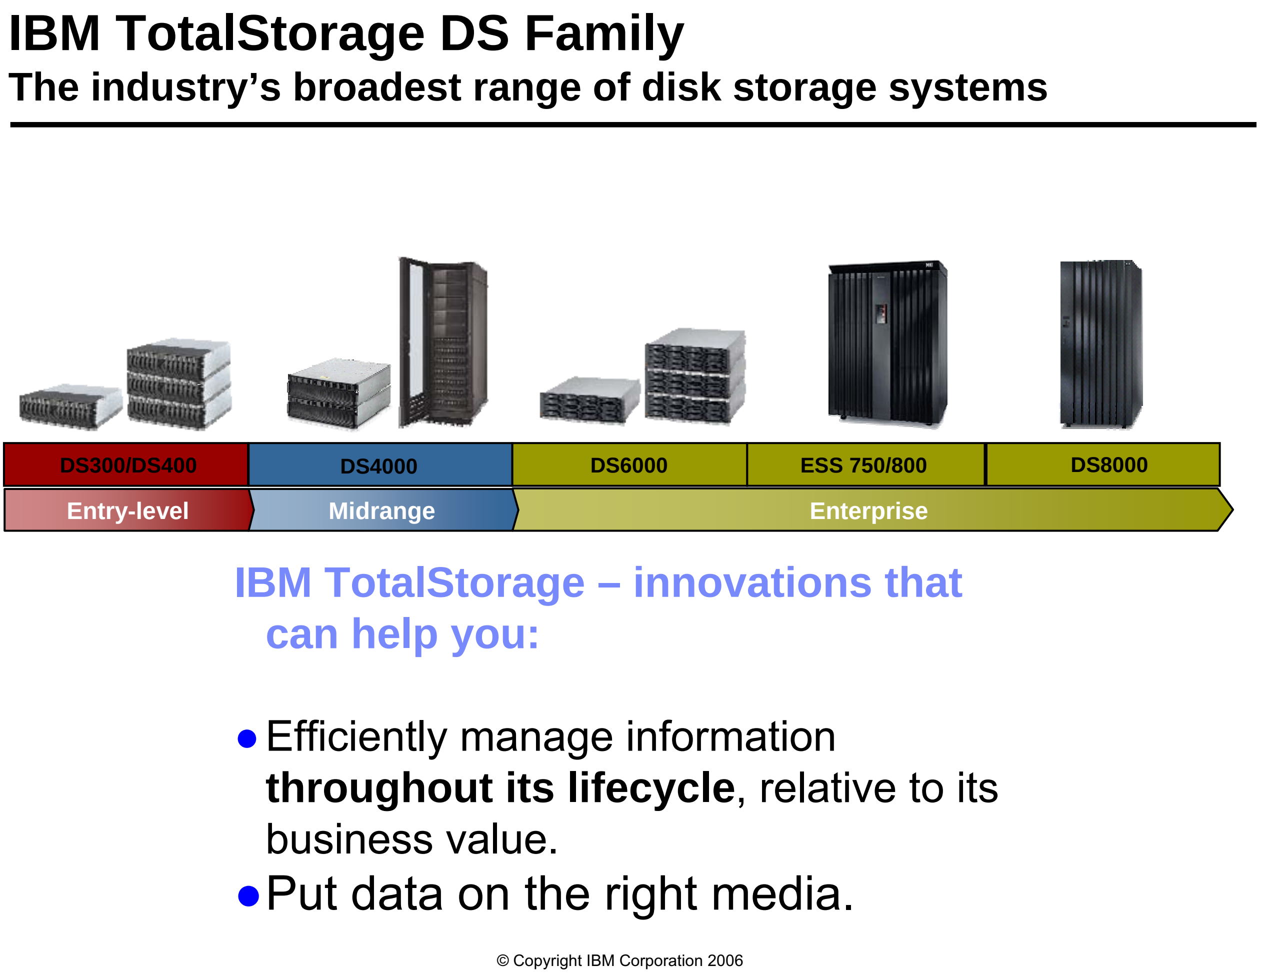

IBM TotalStorage DS Family

The industry’s broadest range of disk storage systems

IBM TotalStorage – innovations that can help you:

- Efficiently manage information **throughout its lifecycle**, relative to its business value.
- Put data on the right media.

Figure 6-17. IBM TotalStorage DS Family The industry's broadest range of disk storage systems AM711.0

Notes:

DS4000 Series

This table represents general positioning and not technical capabilities

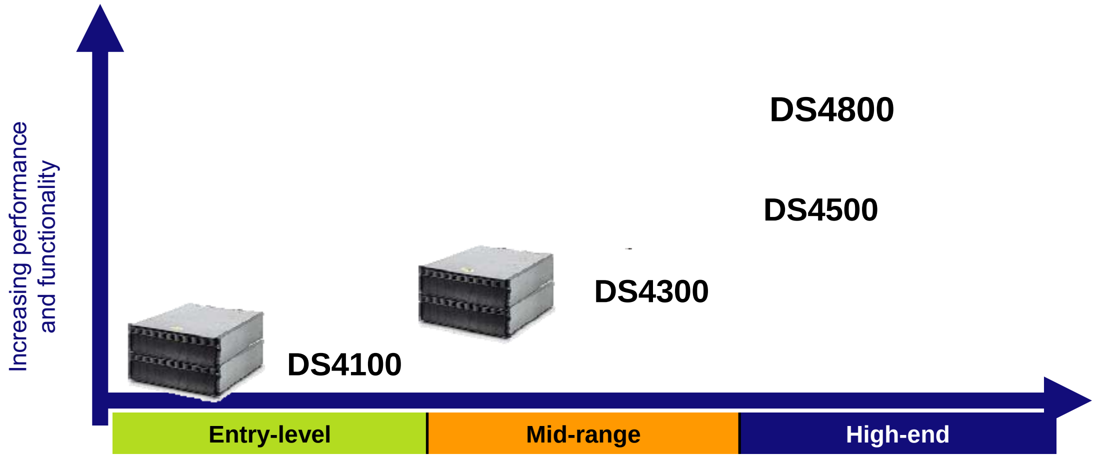

	Entry-level	Mid-range	High-end
Workload	Light workloads	Medium workloads	Heavy workloads
Typical platforms	Windows/Linux direct attach	Homogeneous SAN	Heterogeneous SAN
Applications	SMB operational	SME business-important	SME business-critical
Primary drive type	SATA	FC / SATA	FC
Data replication	FlashCopy	FlashCopy, VolumeCopy, Enhanced Remote Mirroring	FlashCopy, VolumeCopy, Enhanced Remote Mirroring

Figure 6-18. DS4000 Series AM711.0

Notes:

IBM System Storage External Disk Positioning

- Choose IBM TotalStorage DS4000 storage if the client needs:
  - Modular storage starting at lower entry point
  - Great price/performance
  - No zSeries or iSeries support
  - Specific mid-range disk functions/features

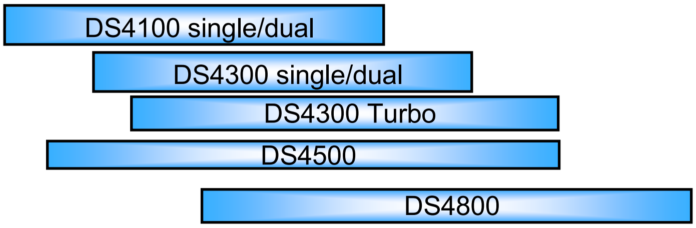

DS4100 single/dual

DS4300 single/dual

DS4300 Turbo

DS4500

DS4800

Figure 6-19. IBM System Storage External Disk Positioning AM711.0

Notes:

This chart provides some perspective on the positioning of the IBM Disk Storage products.

DS6000 Series

- 3U package for controllers and disk expansion units
- Open systems and mainframe support
- Advanced functions interoperate with DS8000, ESS
- High storage density footprint
  - 16 drives per 3U package, including controller
- Up to 7 expansion units
  - Scalable to 128 disk drives; 38.4 TB
- Four year warranty on both hardware and software

DS6800 Controller

Enterprise-class reliability and host attachment features now available in an incredibly small, modular, affordable package

Figure 6-20. DS6000 Series AM711.0

Notes:

The picture is a control unit, but an expansion unit looks exactly the same from the front. Highest density packaging in the industry. 2 key points:

Open and mainframe host attach.

This is an enterprise class box in a small package at very affordable prices.

Do not let the small size fool you, it is a highly available, high performance enterprise class box

There is no room for TotalStorage on the front, it is so small!

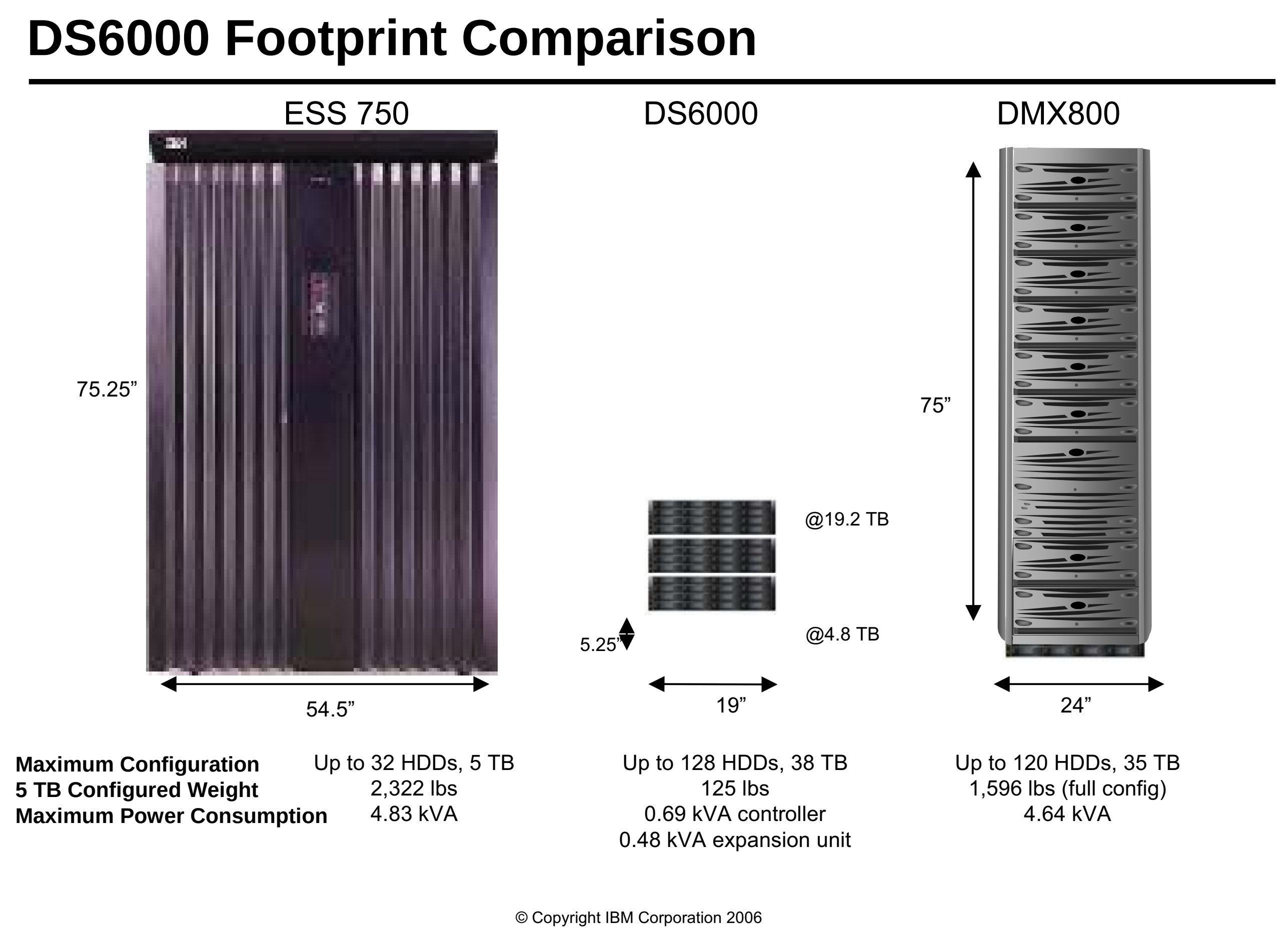

DS6000 Footprint Comparison

	ESS 750	DS6000	DMX800
Maximum Configuration	Up to 32 HDDs, 5 TB	Up to 128 HDDs, 38 TB	Up to 120 HDDs, 35 TB
5 TB Configured Weight	2,322 lbs	125 lbs	1,596 lbs (full config)
Maximum Power Consumption	4.83 kVA	0.69 kVA controller 0.48 kVA expansion unit	4.64 kVA

Figure 6-21. DS6000 Footprint Comparison AM711.0

Notes:

This is self-explanatory. The pictures are drawn to scale.

ESS 750/800

- ESS 750
  - Scales from 1.1 TB to 4.6 TB physical capacity
  - Field upgradeable to ESS 800
- ESS 800
  - Up to 55.9 TB physical capacity
  - Offers high availability and resiliency to support enterprise mission-critical business applications

Figure 6-22. ESS 750/800 AM711.0

Notes:

Designed to provide outstanding price/performance

Scales from 1.1TB up to 4.6TB of physical capacity

Offers high availability with redundant hardware to support mission-critical applications

Supports business continuity with a suite of advanced copy functions designed to protect up-time during planned and unplanned outages

Attaches to a wide-variety of hosts, including all IBM platforms (zSeries®, pSeries®, iSeries™, xSeries®) and a large number of different UNIX®, Linux, and Intel™ processor-based servers

Designed to help lower total costs through efficient, centralized management features

Designed to offer excellent investment protection with field upgradeability to an ESS Model 800, which can scale to over 55TB of physical capacity

ESS EOM and New ESS Offering

- ESS 750/800 EOM
  - All geographies except Asia Pacific
  - Effective April 28, 2006
- Refurbished ESS 800
  - Three year warranty included
  - Effective April 7, 2006
  - MES upgrade capabilities still available

Figure 6-23. ESS EOM and New ESS Offering AM711.0

Notes:

DS8000

- More cost effective scaling made possible by POWER5 technology
- Lower costs through IBM’s virtualization engine technology
- Exponential capacity growth – up to 192 TB capacity
- Dramatic cost reductions available
- Cost effective business continuity capabilities

Figure 6-24. DS8000

Notes:

More cost effective scaling made possible by POWER5TM technology

Expandable configurations - model-to-model field upgrade capability – linear scalability

Up to 256 GB cache and all new cache performance optimization algorithms promote balanced performance up to 6X that of the IBM TotalStorage Enterprise Storage Server® (ESS) Model 800

Switched fiber channel disk technology

Modular design to accommodate future enhancements

Lower costs through IBM's Virtualization Engine Technology

Dual storage system LPARs (logical partitions) offer two independent storage subsystems within a single DS8000 allowing lower costs/mb, administrative control, and application quality of service versus previous alternatives

Additional capabilities planned for future releases

Exponential Capacity Growth – Up to 192TB capacity

IBM System Storage External Disk Positioning

- Choose Enterprise class IBM System Storage if the client needs:
  - – Highest levels of performance and scalability
  - – Storage for zSeries or iSeries
  - – Common advanced copy services across mid-range and Enterprise
  - – Highest serviceability/availability

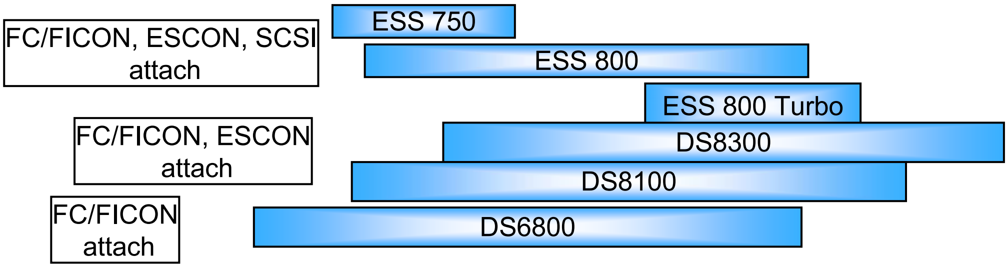

FC/FICON, ESCON, SCSI attach

ESS 750

ESS 800

ESS 800 Turbo

FC/FICON, ESCON attach

DS8300

DS8100

FC/FICON attach

DS6800

Figure 6-25. IBM System Storage External Disk Positioning AM711.0

Notes:

This chart provides some perspective on the positioning of the IBM Disk Storage products.

SAN Switches from SMB to Enterprise

- SAN switch vendors:
  - IBM TotalStorage SAN b-type switch family (Brocade)
  - IBM TotalStorage SAN m-type switch family (McDATA)
  - Cisco MDS 9000 Multilayer Switch Family

BROCADE
The intelligent platform for networking storage

McDATA®

Cisco Systems®

Figure 6-26. SAN Switches from SMB to Enterprise AM711.0

Notes:

SAN Volume Controller Storage Engine

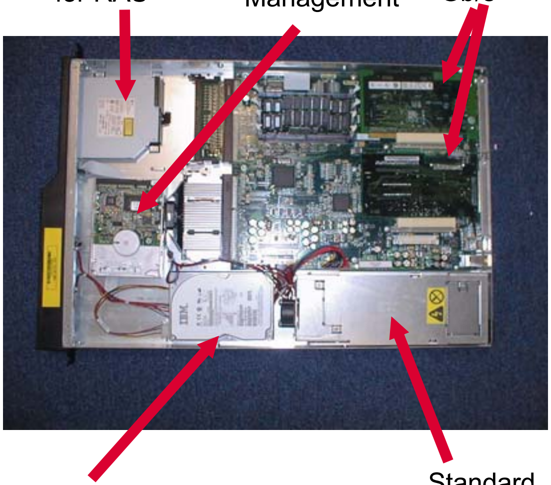

Special Front Panel for RAS

Endeavour High Availability Management

Agilent DX2 Dual 2 Gb/s

Backup and Boot IDE disk

Standard Power supply

SVC Dual Node

IBM Designed Solution

- Linux OS
- Compass distributed cache-clustering SW
- Point in Time Copy
- Peer-to-Peer Remote Copy
- Data Migration Services
- Java-based GUI configurator and CLI
- Writes mirrored to both nodes
- Extensive Disk Storage Attachment

Figure 6-27. SAN Volume Controller Storage Engine AM711.0

Notes:

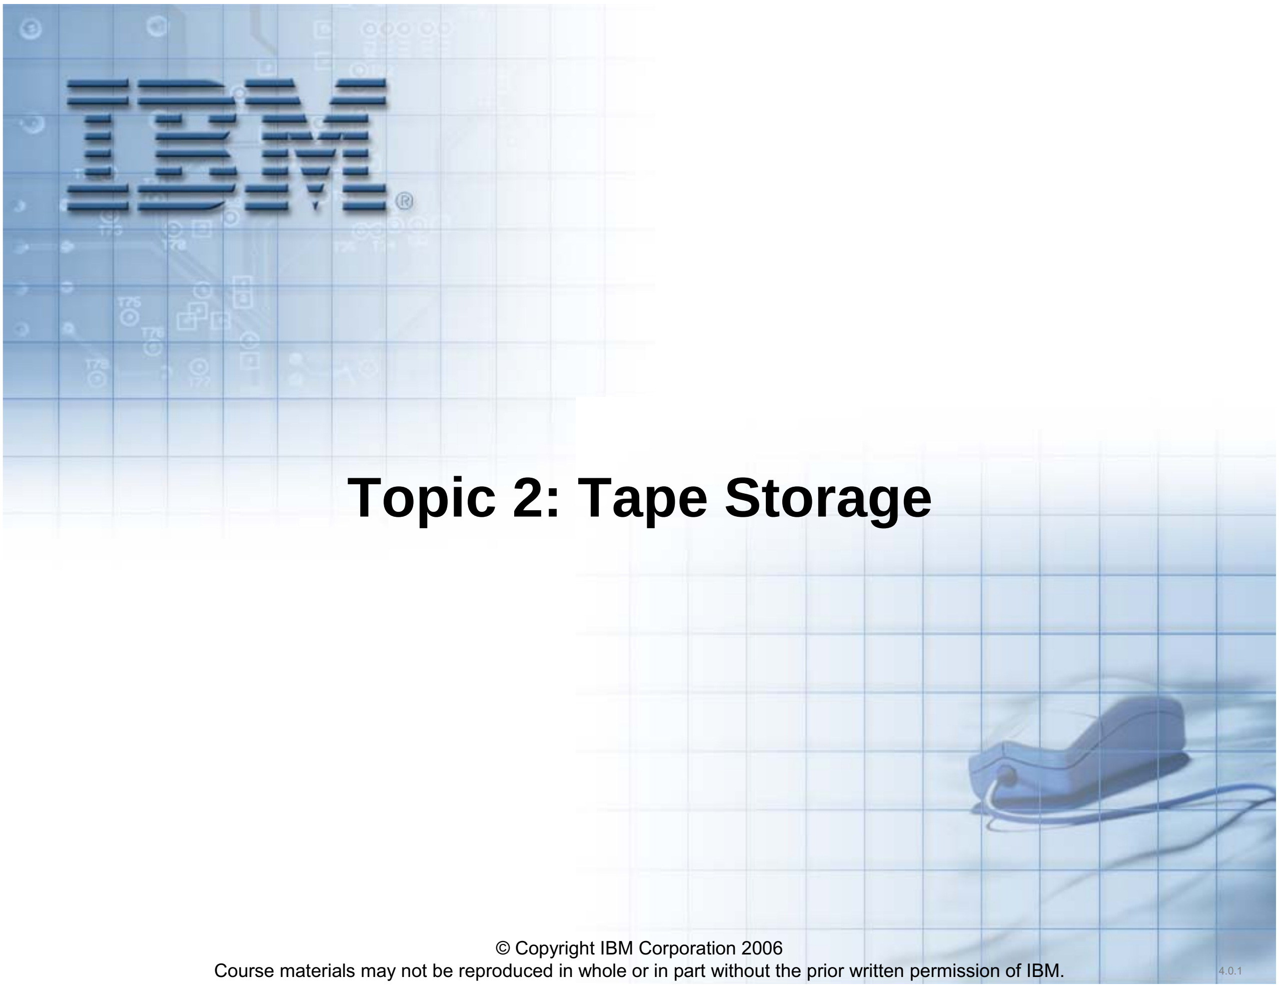

Topic 2: Tape Storage

4.0.1

Figure 6-28. Topic 2: Tape Storage AM711.0

Notes:

IBM TotalStorage LTO Tape Family

Solutions for data archive, backup and disaster recovery

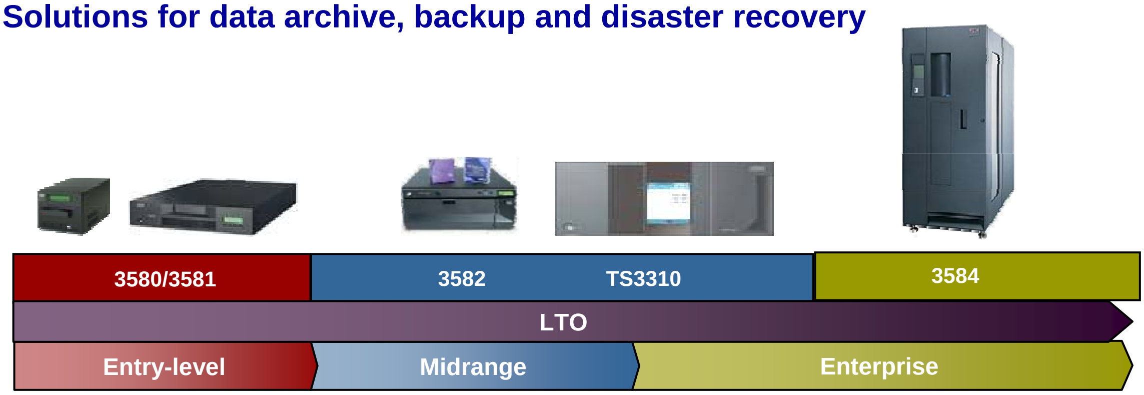
3580/3581
3582
TS3310
3584
LTO
Entry-level
Midrange
Enterprise

IBM TotalStorage – innovations that can help you:

- Efficiently manage information **throughout its lifecycle**, relative to its business value.
  - Putting Data on the Right Media

Figure 6-29. IBM TotalStorage LTO Tape Family AM711.0

Notes:

LTO Tape Drive Roadmap

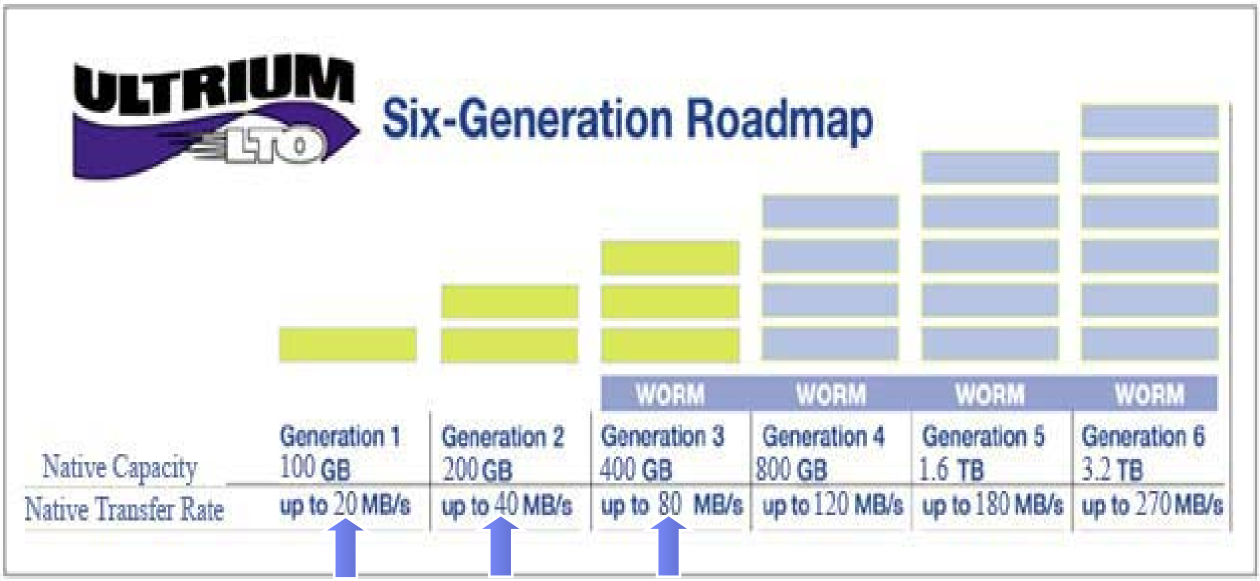

In-the-market LTO Generations

Product Road Map: These statements represent IBM's current intent, are subject to change or withdrawal, and represent only goals and objectives.

Figure 6-30. LTO Tape Drive Roadmap AM711.0

Notes:

IBM LTO Ultrium Offerings

IBM LTO Features

- **Proven:** Over 500,000 IBM LTO drives shipped
- **Open Standard:** Basis for all models
- **High Capacity** 400 GB Native Capacity
- **High Performance:** 80 MB/s Native (up to 575 GB/hr per drive)

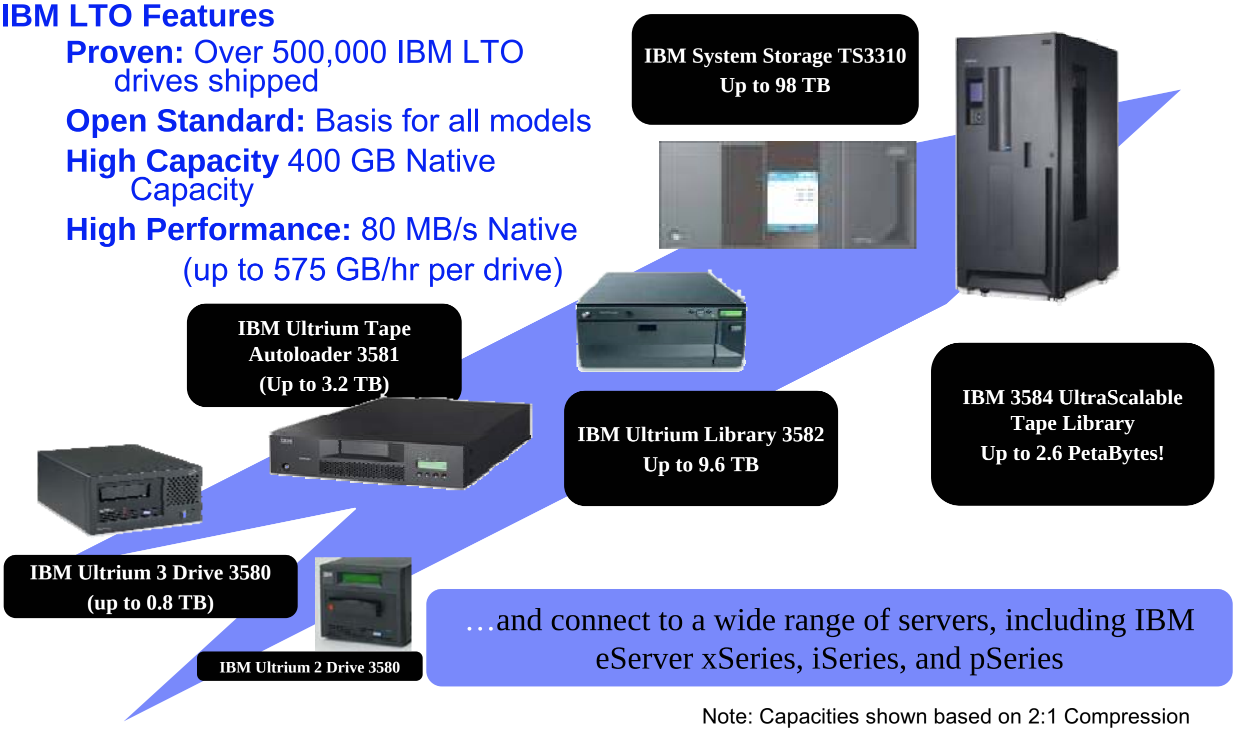

IBM System Storage TS3310
Up to 98 TB

IBM Ultrium Tape Autoloader 3581
(Up to 3.2 TB)

IBM Ultrium Library 3582
Up to 9.6 TB

IBM 3584 UltraScalable Tape Library
Up to 2.6 PetaBytes!

IBM Ultrium 3 Drive 3580
(up to 0.8 TB)

IBM Ultrium 2 Drive 3580

…and connect to a wide range of servers, including IBM eServer xSeries, iSeries, and pSeries

Note: Capacities shown based on 2:1 Compression

Figure 6-31. IBM LTO Ultrium Offerings AM711.0

Notes:

TS3310 Tape Library

- Modular Library
  - 12 TB up to 49 TB storage slot native capacity (30 to 122 LTO storage slots)
  - One to six LTO3 FC/SCSI hot-swappable tape drives (up to 480 MB/sec native performance)
- Supports
  - Logical Partitioning and Native SMI-S
  - Optional features
    - Capacity on Demand
    - Rack Mounting
    - Path Failover (Data and Control)
    - Redundant power
- Host Attachment
  - IBM eServers (except IBM System z9)
  - Selected HP and Sun Microsystems servers
  - Servers running select versions of Microsoft Windows and Linux operating systems

Figure 6-32. TS3310 Tape Library AM711.0

Notes:

Basic to all models is logical library partitioning, built upon IBM's patented Multipath architecture. This new series of libraries also supports Native SMI-S. Optionally, users can exercise “On Demand” capacity expansion of the stacked unit shown in this slide. A rack mounting kit as well as a second power supply are also being offered as optional features, orderable either by feature code at time of the initial order or later as a field upgrade. Data Path and Control Path failover, first introduced on the IBM 3584 tape library, can also optionally be added to the TS3310. For convenience, both are ordered in one single feature code option.

As one would expect from IBM, the TS3310 tape library attaches to one of the broadest varieties of host systems in the world, from our own eServers (with the exception of zSeries or z9 as it is being re-branded) to Sun, HP and a very wide variety of other servers.

Statement of Direction

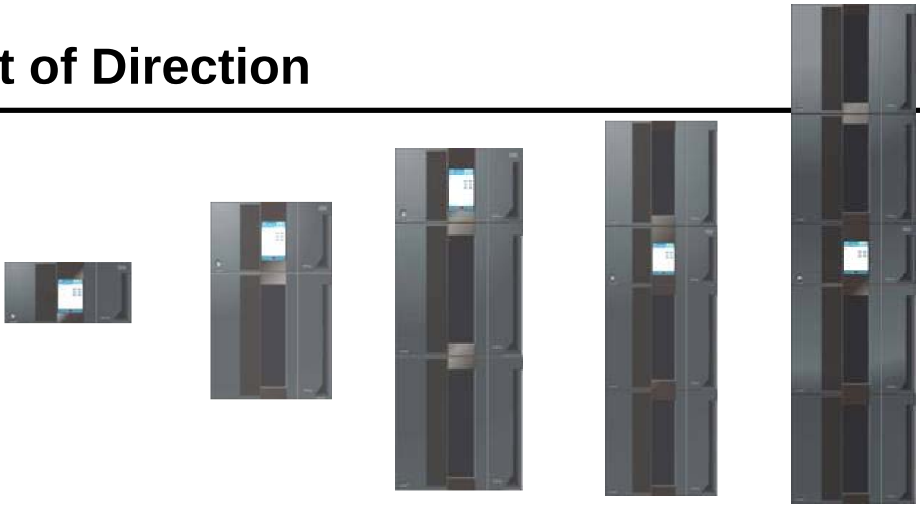

Height	5u	14u	23u	32u	41u
Standalone / Rack mount	Yes / Yes	Yes / Yes	No / Yes	No / Yes	No / Yes
Height in EIA units (1u=1.75", 44.5mm)	5u	14u	23u	32u	41u
Native Capacity (TB) – without I/O	12	49	86	122	158
I/O (slots)	6 (fixed)	up to 18	up to 30	up to 42	up to 54
Storage (slots)	up to 30	up to 122	up to 214	up to 306	up to 398
Drives	up to 2	up to 6	up to 10	up to 14	up to 18

Product Road Map: These statements represent IBM's current intent, are subject to change or withdrawal, and represent only goals and objectives.

Figure 6-33. Statement of Direction AM711.0

Notes:

In this slide we are showing the complete future roadmap for the TS3310. Rollout of these additional models as well as the field upgradeability of previously installed, smaller models is intended to provide investment protection and option up-sell capability to the install base for a number of years.

The second reason for showing this chart is to help you strategize on the selling of the TS3310 versus the 3584. Keep in mind that while the TS3310 will eventually scale-up to 404 cartridges and 18 drives, the 3584 gives customers well over 10 times that capacity. So, for customers who begin looking at future models of the TS3310 and potentially envision maxing-out a TS3310 future configuration, you should definitely switch to the discussion to the benefits of the 3584.

The break point for selling the 3584 versus TS3310 is around 100TBs and 10 drives (in round #s).

IBM TotalStorage Enterprise Tape Family

- Efficiently manage information **throughout its lifecycle**, relative to its business value
  - Putting Data on the Right Media

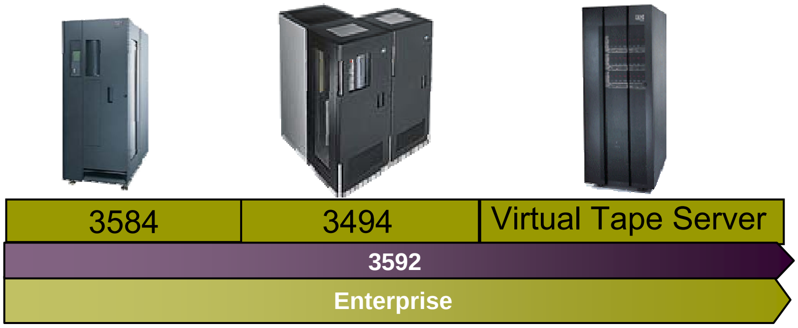

3584

3494

Virtual Tape Server

3592

Enterprise

Figure 6-34. IBM TotalStorage Enterprise Tape Family AM711.0

Notes:

3584 Tape Library

- Modular, scalable design features multipath architecture
  - 1 – 16 frames; up to 12 drives per frame (192 total)
  - 5.6 petabytes of data total (compressed)

Supporting
LTO Ultrium 1
LTO Ultrium 2
LTO Ultrium 3
IBM 3592 tape drives and media within the same library

Figure 6-35. 3584 Tape Library AM711.0

Notes:

3590 Tape Drive

- 3590 tape drive
  - Enhanced tape sharing with cross enterprise attachment
  - High performance for increased productivity
  - High capacity for reduced space consumption and reduced tape cartridge handling
  - Improved data integrity and reliability for mission critical data
  - 3590 is supported for native attachment to Ultra SCSI, fibre channel, ESCON and FICON servers
  - Features an uncompressed drive data rate of up to 14 MB/second
  - Uses 3590 extended length tape cartridge for up to 180 GB of 3:1 compressed data per cartridge

Figure 6-36. 3590 Tape Drive AM711.0

Notes:

3592 Tape Drive

- Supports IBM and selected open system platforms
- Supported on existing IBM and StorageTek automation
- Offers native data transfer rate of up to 40 MB/sec
- Supports up to 900 GB on a cartridge
- Supports fast access to data with 60 GB cartridge
- Supports consolidation to help reduce tape TCO

Figure 6-37. 3592 Tape Drive AM711.0

Notes:

IBM TotalStorage Tape Library/Server Offerings

- 3494 tape library
  - Designed to provide reliable, scalable tape automation
  - Provides multiplatform connectivity
  - Supports 3592 rewritable and WORM cartridges
  - Supports multiple IBM tape drive models
  - Supports IBM virtual tape

- 3494 virtual tape server
  - Can help reduce real tape mounts, by using a tape volume cache (TVC)
  - Can reduce physical tape cartridges required because of higher utilization of cartridge capacity
  - Can help reduce operating costs such as power, maintenance, operations and support staff
  - Can help reduce floor space required to support the tape process, as a result of fewer physical resources
  - Can help to improve performance due to the elimination of most of the physical movement of tape
  - Can be upgraded to peer-to-peer configuration to support business continuance

Figure 6-38. IBM TotalStorage Tape Library/Server Offerings

AM711.0

Notes:

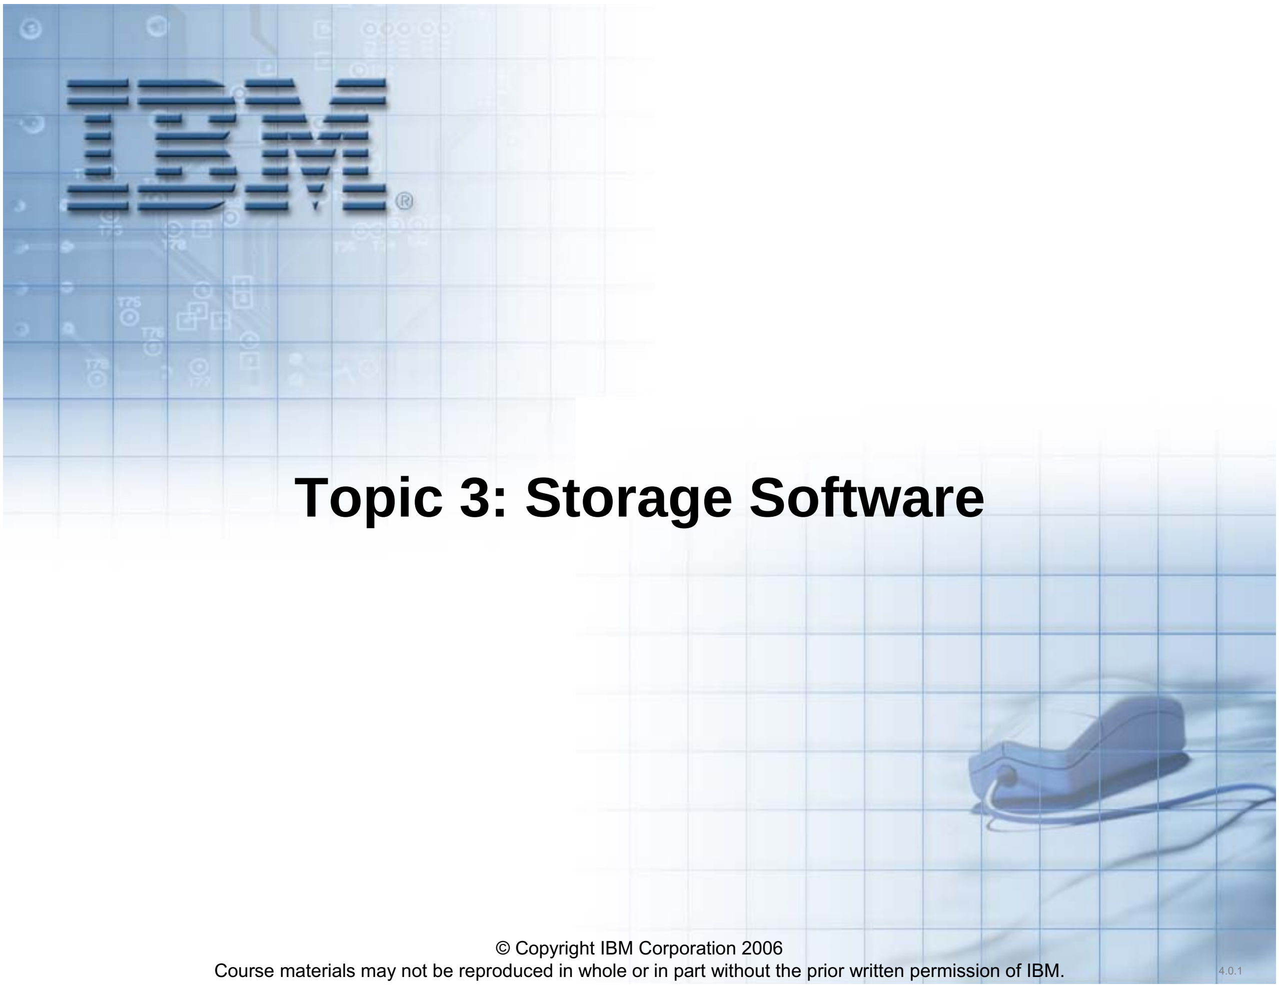

Topic 3: Storage Software

4.0.1

Figure 6-39. Topic 3: Storage Software AM711.0

Notes:

IBM TotalStorage Open Software Storage Family

- Broad portfolio
- Modular integrated and interoperable with those of many other vendors
- Based on open industry-supported standards.
- Includes Virtualization Engine Suite for Storage
  - SAN Volume Controller SVC
  - IBM TotalStorage SAN File System
  - The IBM TotalStorage Productivity Center
- Tivoli Storage Manager family

Figure 6-40. IBM TotalStorage Open Software Storage Family AM711.0

Notes:

Disk Hardware Virtualization

SANs Today **Block Virtualization**

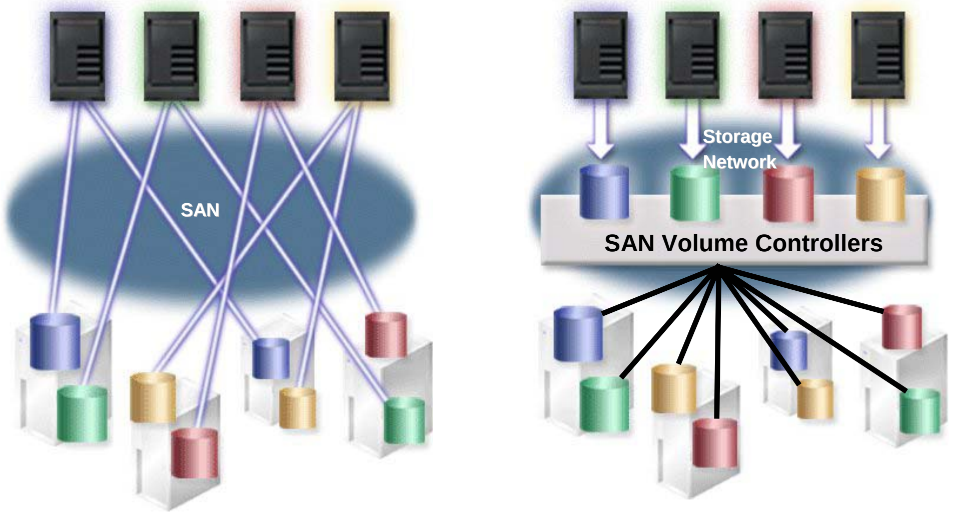

Servers are mapped to specific physical disks
that is, physical mapping

Servers are mapped to a virtual disk
that is, logical mapping

Figure 6-41. Disk Hardware Virtualization AM711.0

Notes:

SAN Volume Controller (SVC) V3.1

- Combine capacity into a single storage pool
- Manage storage pool from a central point
- Change physical storage systems with minimal or no impact to the applications
- Increase scalability
- Reduce downtime
- Increase
  - Storage capacity utilization
  - Uptime
  - Administrator productivity and efficiency
- Single set of advanced copy services across multiple heterogeneous storage systems

Block Virtualization

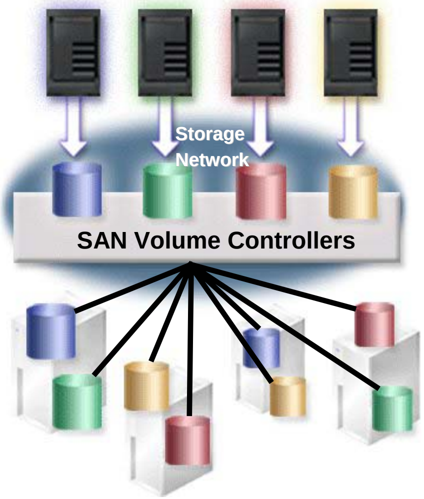

Servers are mapped to a virtual disk
that is, logical mapping

Figure 6-42. SAN Volume Controller (SVC) V3.1 AM711.0

Notes:

SAN Volume Controller (SVC) for Cisco MDS 9000

- The benefits of the IBM TotalStorage SAN Volume Controller embedded in a highly available SAN switch platform
- The SAN Volume Controller for Cisco MDS 9000 combines IBM's storage virtualization software with Caching Services Modules from the Cisco MDS 9000 family of switches and directors to provide these benefits.

Figure 6-43. SAN Volume Controller (SVC) for Cisco MDS 9000 AM711.0

Notes:

SAN File System (SFS)

- Virtualize and consolidate file systems across Microsoft Windows, UNIX, and Linux servers
- Provide a common pool of storage that can be shared across servers
- Support multiple tiers of storage
- Provide policy-based allocation of space at a file level
- Policy-based migration of files between storage pools

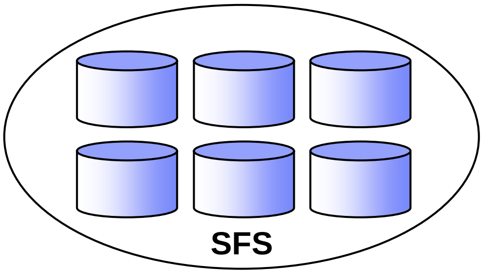

Figure 6-44. SAN File System (SFS) AM711.0

Notes:

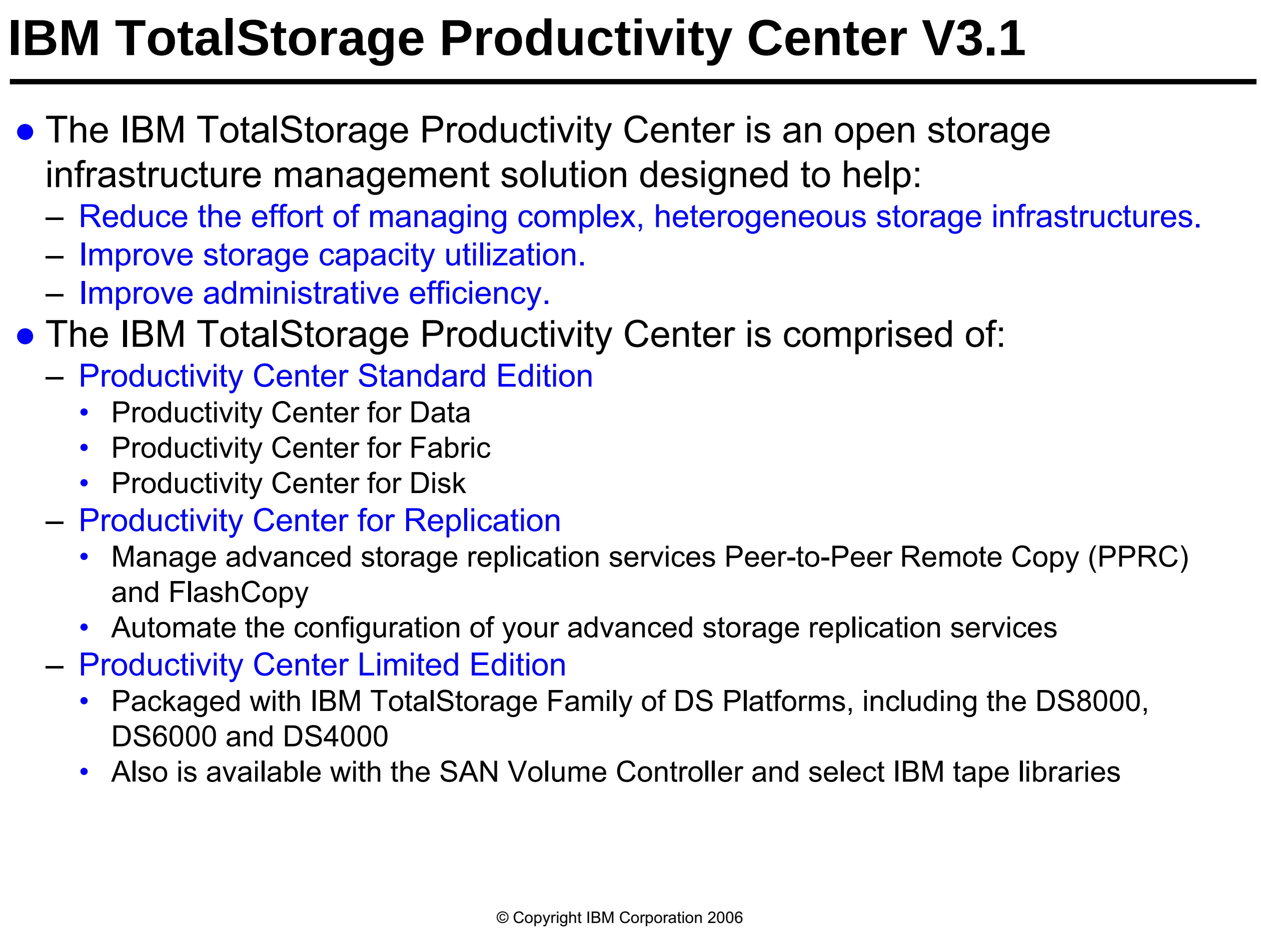

IBM TotalStorage Productivity Center V3.1

- The IBM TotalStorage Productivity Center is an open storage infrastructure management solution designed to help:
  - Reduce the effort of managing complex, heterogeneous storage infrastructures.
  - Improve storage capacity utilization.
  - Improve administrative efficiency.
- The IBM TotalStorage Productivity Center is comprised of:
  - Productivity Center Standard Edition
    - Productivity Center for Data
    - Productivity Center for Fabric
    - Productivity Center for Disk
  - Productivity Center for Replication
    - Manage advanced storage replication services Peer-to-Peer Remote Copy (PPRC) and FlashCopy
    - Automate the configuration of your advanced storage replication services
  - Productivity Center Limited Edition
    - Packaged with IBM TotalStorage Family of DS Platforms, including the DS8000, DS6000 and DS4000
    - Also is available with the SAN Volume Controller and select IBM tape libraries

Figure 6-45. IBM TotalStorage Productivity Center V3.1 AM711.0

Notes:

Tivoli Storage Manager

- **IBM Tivoli Storage Manager family of offerings include:**
  - IBM Tivoli Storage Manager for Application Servers
  - IBM Tivoli Storage Manager for Databases
  - IBM Tivoli Storage Manager for Enterprise Resource Planning
  - IBM Tivoli Storage Manager for Hardware
  - IBM Tivoli Storage Manager for Mail
  - IBM Tivoli Storage Manager for System Backup and Recovery
- **Optional Additions to Tivoli Storage Manager:**
  - **Hierarchical Storage Management Products**
    - IBM Tivoli Storage Manager for Space Management
    - IBM Tivoli Storage Manager HSM for Windows
  - **Complementary Products:**
    - IBM Tivoli Continuous Data Protection for Files
    - IBM Tivoli Storage Manager for Storage Area Networks

Tivoli. software

Figure 6-46. Tivoli Storage Manager AM711.0

Notes:

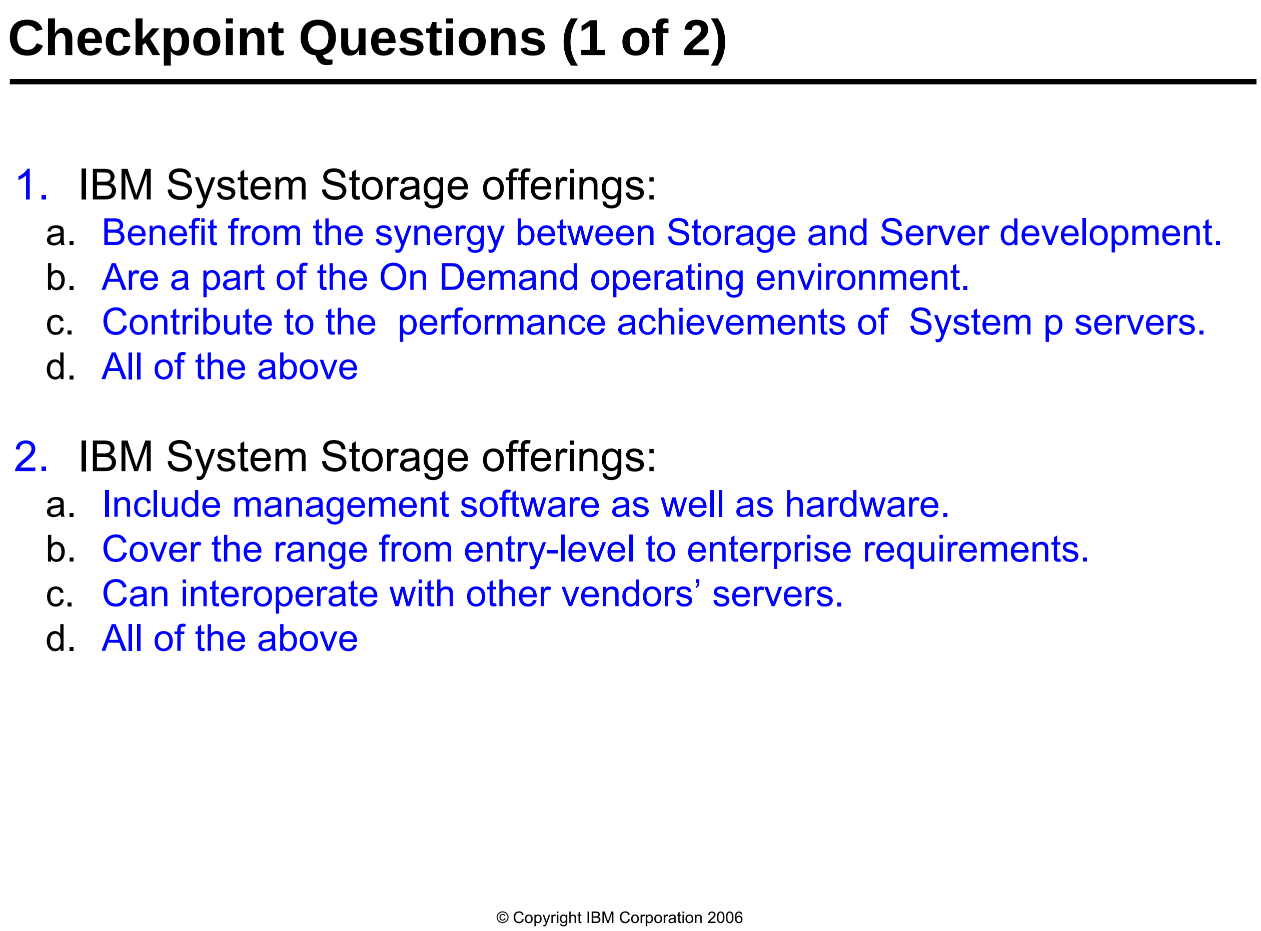

Checkpoint Questions (1 of 2)

1. IBM System Storage offerings:
   a. Benefit from the synergy between Storage and Server development.
   b. Are a part of the On Demand operating environment.
   c. Contribute to the performance achievements of System p servers.
   d. All of the above

2. IBM System Storage offerings:
   a. Include management software as well as hardware.
   b. Cover the range from entry-level to enterprise requirements.
   c. Can interoperate with other vendors’ servers.
   d. All of the above

Figure 6-47. Checkpoint Questions (1 of 2) AM711.0

Notes:

Checkpoint Questions (2 of 2)

3. The SAN Volume Controller:
   a. Combines storage capacity into a single storage pool.
   b. Allows management of the storage pool from a central point.
   c. Reduces down time.
   d. All of the above

4. True or False. SAN and NAS are simply different names for the same storage concept.

5. LTO Ultrium 3 tapes have a compressed storage capacity of ______.
   a. 100 Gigabytes
   b. 200 Gigabytes
   c. 800 Gigabytes
   d. 1000 Gigabytes

Figure 6-48. Checkpoint Questions (2 of 2)

Notes:

Unit Summary

You should now be able to:

- Position IBM System Storage disk products.
- Position IBM System Storage tape products.
- Discuss IBM System Storage software solutions.

Figure 6-49. Unit Summary AM711.0

Notes:

Any Questions?

Figure 6-50. Any Questions? AM711.0

Notes:

Unit 7. Software

What This Unit Is About

This unit describes the software products relevant to System p solution design

What You Should Be Able to Do

After completing this unit, you should be able to:

- Discuss the benefits of AIX
- Describe the benefits of Linux on Power
- Discuss IBM software products and solutions

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

After completing this unit, you should be able to:

- Discuss the benefits of AIX.
- Describe the benefits of Linux on Power.
- Discuss IBM software products and solutions.

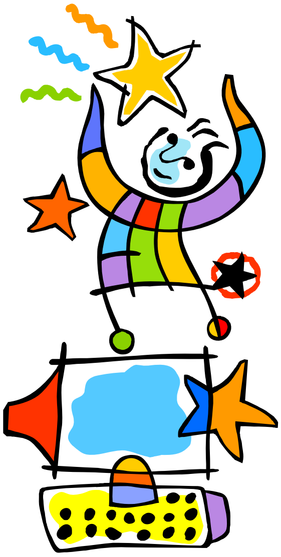

Figure 7-1. Unit Objectives AM711.0

Notes:

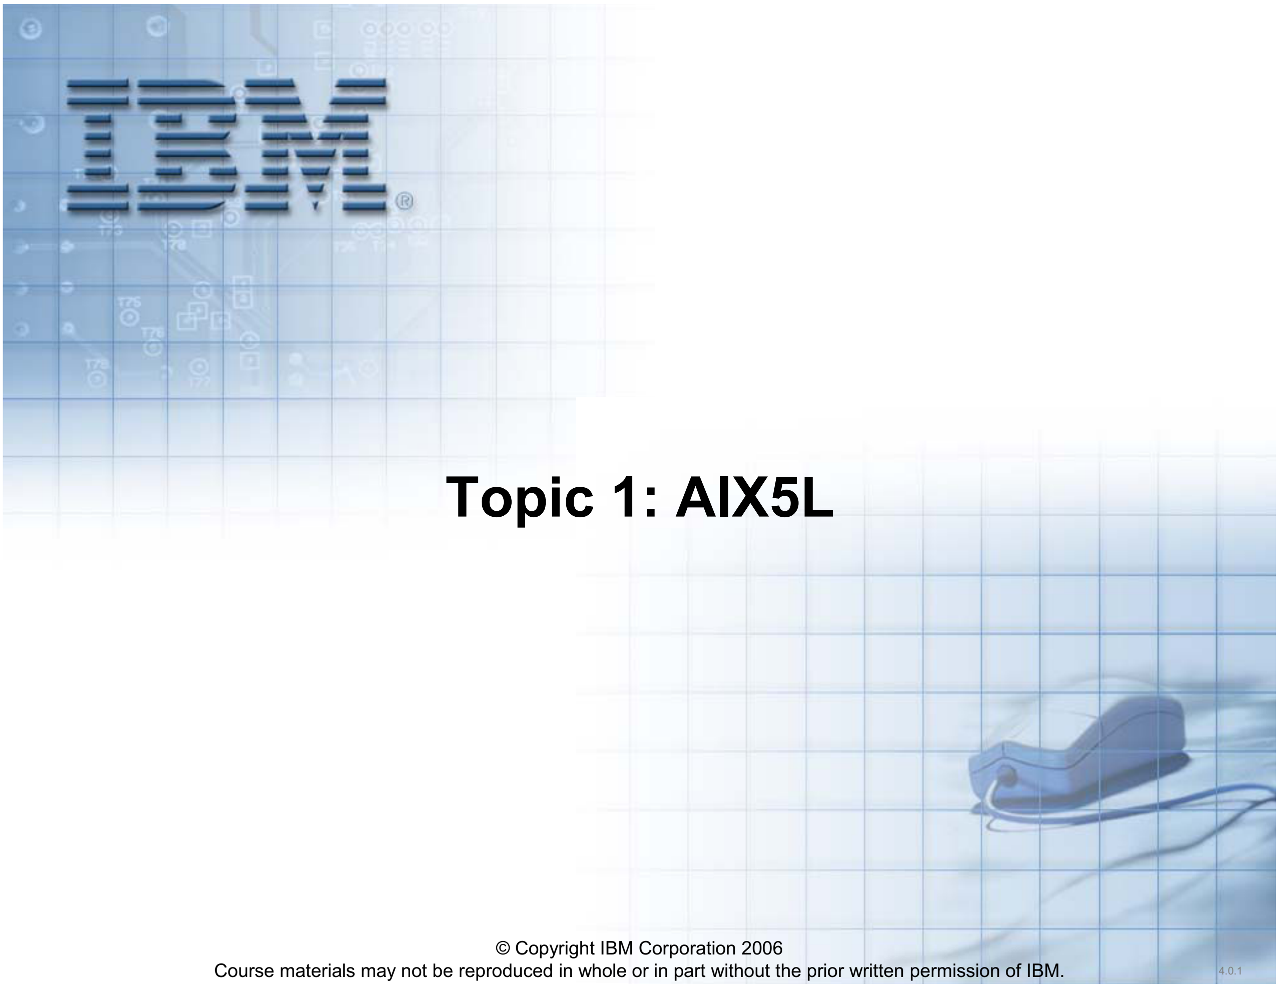

Topic 1: AIX5L

4.0.1

Figure 7-2. Topic 1: AIX5L AM711.0

Notes:

Topic Objectives

- On completion of this topic you should be able to:
  - Describe the current UNIX market.
  - Position AIX in that market.
  - Describe to clients the benefits of AIX versus other UNIX operating systems.
  - Describe the new features included in AIX5L V5.3 TL5.

Figure 7-3. Topic Objectives AM711.0

Notes:

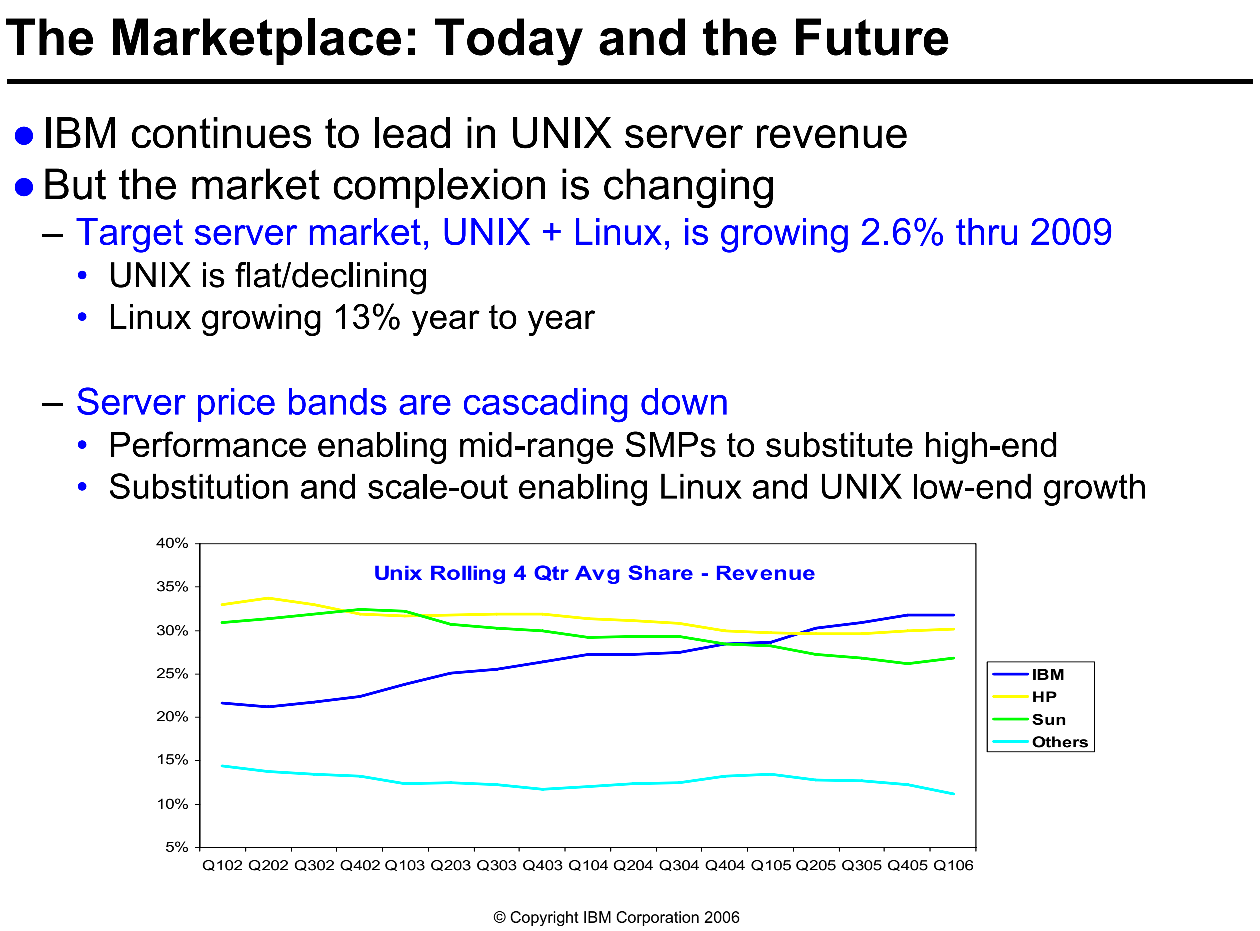

The Marketplace: Today and the Future

- IBM continues to lead in UNIX server revenue
- But the market complexion is changing
  - Target server market, UNIX + Linux, is growing 2.6% thru 2009
    - UNIX is flat/declining
    - Linux growing 13% year to year
  - Server price bands are cascading down
    - Performance enabling mid-range SMPs to substitute high-end
    - Substitution and scale-out enabling Linux and UNIX low-end growth

Unix Rolling 4 Qtr Avg Share - Revenue
40%
35%
30%
25%
20%
15%
10%
5%
Q102 Q202 Q302 Q402 Q103 Q203 Q303 Q403 Q104 Q204 Q304 Q404 Q105 Q205 Q305 Q405 Q106
IBM
HP
Sun
Others

Figure 7-4. The Marketplace: Today and the Future AM711.0

Notes:

What Do Clients Want from an Operating System?

- Support for the application they want to deploy
- Excellent reliability and availability
- Ease of administration and system management
- Built-in security
  - Protection
  - Access control
- Flexibility
  - Ability to adjust dynamically
- Low total cost of ownership

Figure 7-5. What Do Clients Want from an Operating System? AM711.0

Notes:

Why AIX5L?

- Mainframe-inspired
- Broad application support
- Compelling performance
- Long term roadmap

Figure 7-6. Why AIX5L? AM711.0

Notes:

Features of AIX5L V3 on Power5

- Flexible Resource Model
- System Scalability
- Security
- Reliability, Availability, Serviceability (RAS)
- Network Performance
- Systems Management
- Storage
- Development and Performance

Figure 7-7. Features of AIX5L V3 on Power5 AM711.0

Notes:

AIX 5L V5.3 Technology Level 5

- New hardware support
- Security enhancements
- Enterprise reliability
- Shared boot Image management
- NIM and install enhancements
- NFS Proxy Caching
- Network Data Administration Facility (NDAF)
- NFSv4 grace period support
- Improved performance tools
- AIX 5L command enhancements
- Networking enhancements
- VIOS V1.3
- Tivoli integration
- Binary compatibility with previous releases

New August 2006

Figure 7-8. AIX 5L™ V5.3 Technology Level 5 AM711.0

Notes:

Checkpoint Questions (1 of 6)

1. Which of the following are true statements about AIX 5L as a market leader among UNIX operating systems?
   a) Mainframe-inspired capabilities
   b) Broad application support
   c) Available on Intel and Power architecture
   d) Statements a and b are true

2. The client is interested in being able to move IT resources dynamically. Which of the following advantages of AIX on Power5 architecture would you emphasize?
   a. Reliability, availability, and serviceability.
   b. Flexible resource model
   c. Network performance
   d. Systems management

Figure 7-9. Checkpoint Questions (1 of 6) AM711.0

Notes:

Checkpoint Questions (2 of 6)

3. Currently the UNIX marketplace is characterized by:
   a. Rapid growth and the emergence of numbers of small vendors.
   b. A rapid decline as Windows operating systems dominate all segments.
   c. Flat growth or small declines with fierce competition for market share
   d. Rapid growth of Solaris

4. Which of the following is feature introduced at AIX5L V5.3 TL5?
   a. JFS2 file system support
   b. Shared boot image management
   c. NFSV3 support
   d. KDE window manager support

Figure 7-10. Checkpoint Questions (2 of 6) AM711.0

Notes:

Topic Summary

- Having completed this topic you should be able to:
  - Describe the current UNIX market.
  - Position AIX in that market.
  - Describe to clients the benefits of AIX versus other UNIX operating systems.
  - Describe the new features included in AIX5L V5.3 TL5.

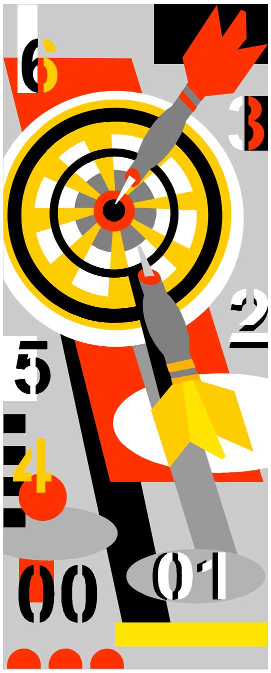

Figure 7-11. Topic Summary AM711.0

Notes:

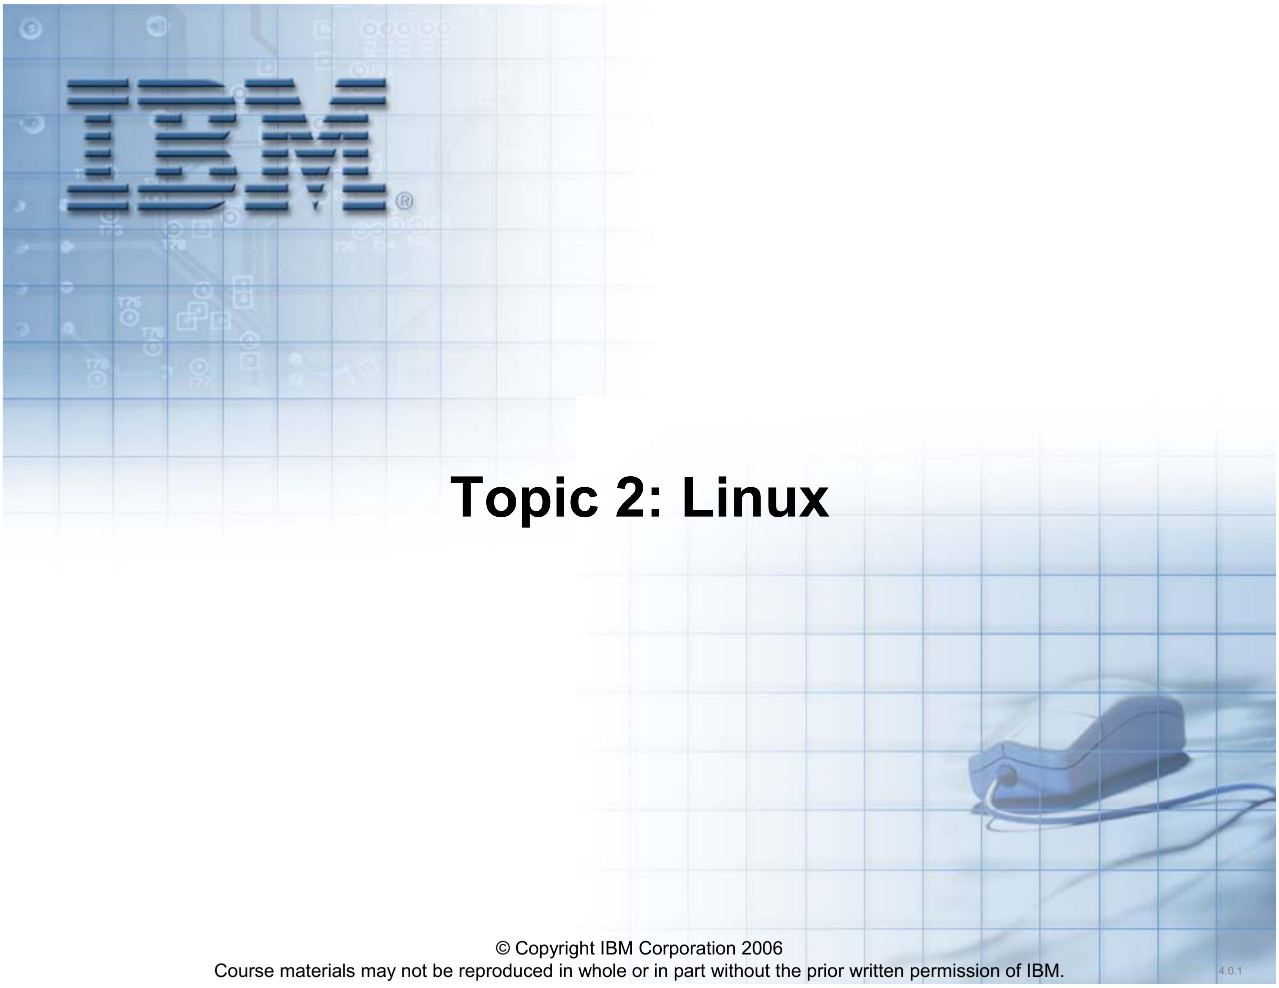

Topic 2: Linux

4.0.1

Figure 7-12. Topic 2: Linux AM711.0

Notes:

Topic Objectives

- After completing this topic you should be able to:
  - Describe the Linux market and position Power architecture systems within that market.
  - Describe the market opportunity created by Linux for IBM System p.
  - Describe the major features of Linux to clients.
  - Describe the Linux on Power benefits versus Linux on Intel.

Figure 7-13. Topic Objectives AM711.0

Notes:

The Operating Systems Battlefield

- Core strength for Unix is in the >$25K space
- Core strength for Windows is in the <$6K space
- The space in between represents the areas to be contested

Price Band	OS->	Windows		Linux		Unix	
		’06 Oppy	’06-’09 CGR	’06 Oppy	’06-’09 CGR	’06 Oppy	’06-’09 CGR
<$6K	Rev	$15.3B	6%	$4.5B	12%	$0.9B	2%
	Units	4,722K	11%	1,454K	18%	308K	8%
$6K-$25K	Rev	$3.9B	8%	$1.5B	13%	$3.0B	-1%
	Units	335K	12%	133K	16%	259K	-1%
>$25K	Rev	$1.2B	15%	$0.7B	21%	$14.3B	-1%
	Units	21K	12%	9K	22%	136K	-3%
Totals	Rev	$20.5B	7%	$6.7B	13%	$18.2B	-1%
	Units	5,078K	11%	1,596K	18%	702K	3%

Source: IDC ’06 updated

RFC Ref: mkt overview.xls

Figure 7-14. The Operating Systems Battlefield

AM711.0

Notes:

UNIX Revenue Share – Rolling 4 Quarters

System p has continually gained share to take and maintain its Unix leadership position

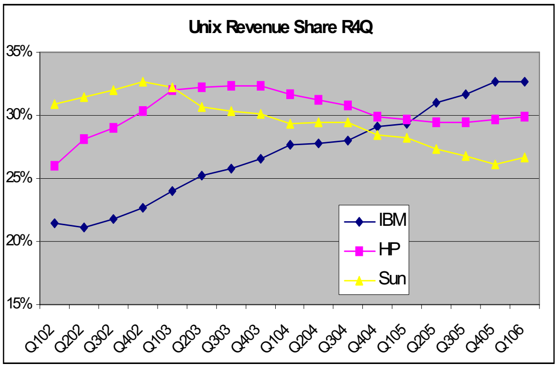

In Linux, IBM has steadily gained share and is closing the gap.

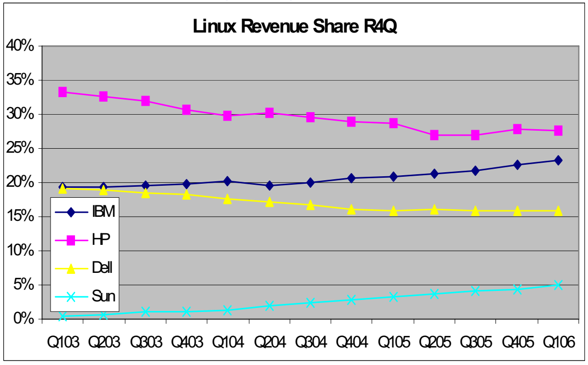

Vast bulk of IBM Linux is Intel x86 based, POWER's contribution is in 10% range; Opteron, 12%; Itanium, 2%

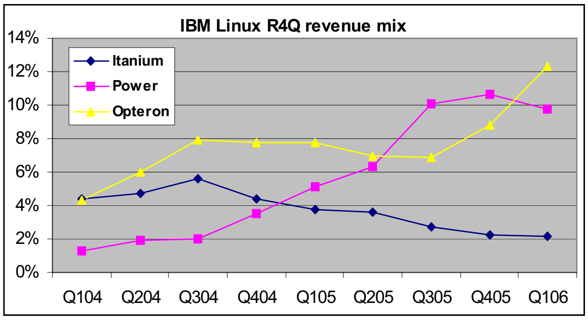

Source: IDC updated

RFC Ref: mkt overview.xls

Figure 7-15. UNIX Revenue Share – Rolling 4 Quarters AM711.0

Notes:

What is Linux?

- **UNIX-like operating system**
- **Developed / tested by the Open Source community**
  - **POWER contributions by IBM LTC**
- **Packaged and shipped by distributors**
  - **Red Hat, Novell SuSE, Debian**
  - **Other regional distributors**

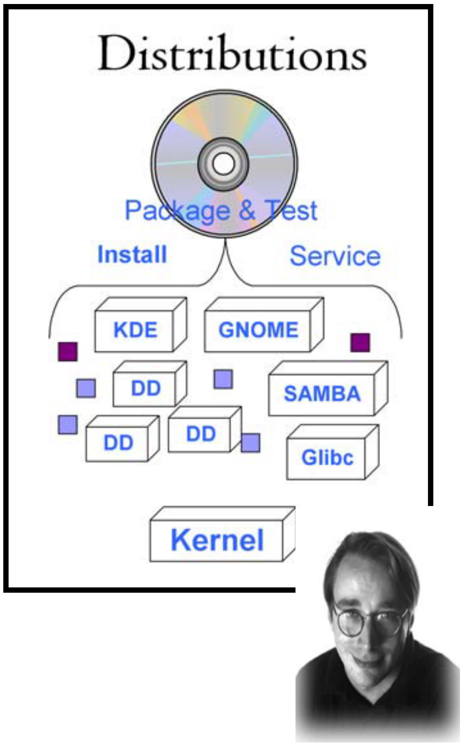

> *"Hello everybody... I'm doing a (free) operating system (just a hobby, won't be big and professional...)."*
>
> Linus Torvalds, creator of Linux, from the first Internet announcement on August 25, 1991. Even he initially underestimated its potential.

Figure 7-16. What is Linux? AM711.0

Notes:

Why Linux on Power?

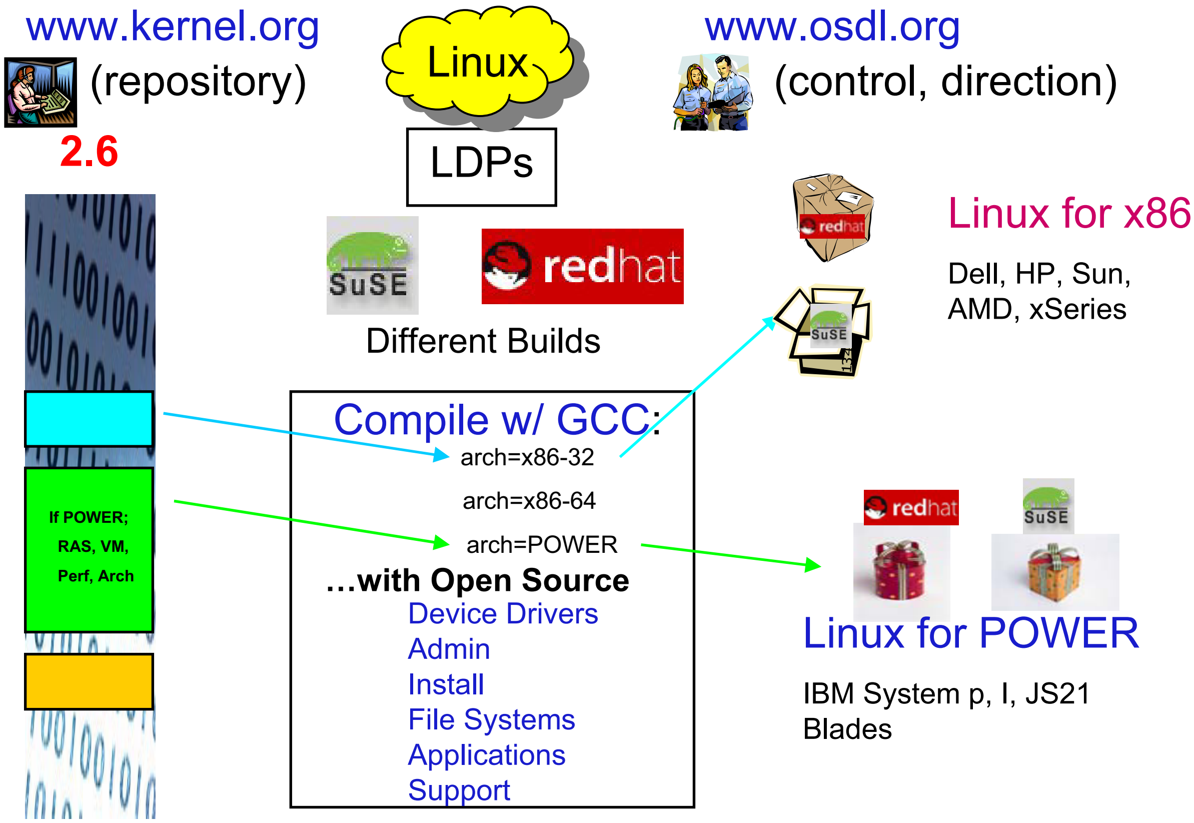

www.kernel.org
(repository)
2.6

Linux
LDPs

www.osdl.org
(control, direction)

Different Builds

Linux for x86
Dell, HP, Sun,
AMD, xSeries

If POWER;
RAS, VM,
Perf, Arch

Compile w/ GCC:
arch=x86-32
arch=x86-64
arch=POWER

…with Open Source
Device Drivers
Admin
Install
File Systems
Applications
Support

Linux for POWER
IBM System p, I, JS21
Blades

Figure 7-17. Why Linux on Power? AM711.0

Notes:

Linux and IBM System p Virtualization

- Award winning combination
- **Linux Journal and IDG World Expo Announce Winners of Product Excellence Awards at LinuxWorld San Francisco**

BEST VIRTUALIZATION SOLUTION

IBM Advanced POWER Virtualization

Figure 7-18. Linux and IBM System p Virtualization AM711.0

Notes:

IBM Linux on POWER Strategy

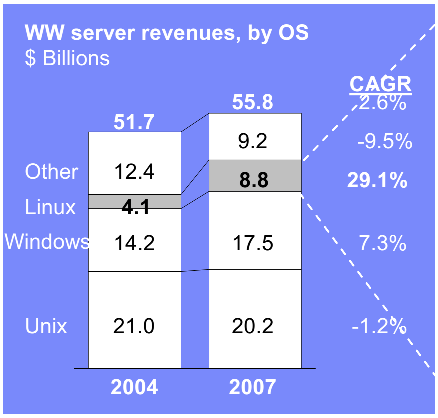

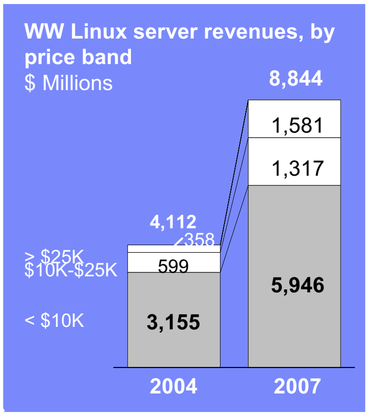

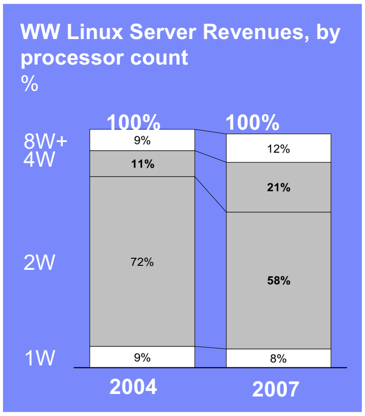

- **Linux on POWER is strategic to IBM**
  - Capitalize on Linux market growth
  - Strengthen low-end position
  - Driver of new footprints (for example, Blades)
  - Provide customers platform choice
  - Stimulate the ecosystem around POWER
- **IBM is investing for success**
  - High investment level within server group
  - New group formed
  - Dedicated Geo teams
  - One IBM support

Evolving from Accommodation to Leadership

Figure 7-19. IBM Linux on POWER Strategy AM711.0

Notes:

POWER Linux Discussion Points with Customers

- People are choosing Linux in your accounts
- Default is Intel-based systems unless you promote:
  - Performance
    - Comparable Linux and AIX workload performance
    - Single POWER architecture for 64-bit or 32-bit high performance
    - Large memory addressability and bandwidth
    - Highly scalable horizontal or vertical growth on Demand
  - Application Capacity
    - Extensive system capacity - memory, CPU, I/O, storage
    - Dedicated, clustered, or partitioned Linux
    - Native Linux partitions to complement AIX
    - Enterprise Class Application Support
    - Up to 160 logical partitions with dynamic LPARs and Virtual I/O capabilities
  - Tuned for Linux
    - Proven System p hardware and on Demand and RAS features extend Linux to new level

Figure 7-20. POWER Linux Discussion Points with Customers AM711.0

Notes:

Sources of Information

- For clients:

http://www-1.ibm.com/linux/

- Latest Announcements
- Linux Express Configs
- Solutions
- Links
  - SuSE
  - RedHat
  - Linux Lens
  - Linux Apps Redbook
- Whitepapers
- Facts and Features

Figure 7-21. Sources of Information AM711.0

Notes:

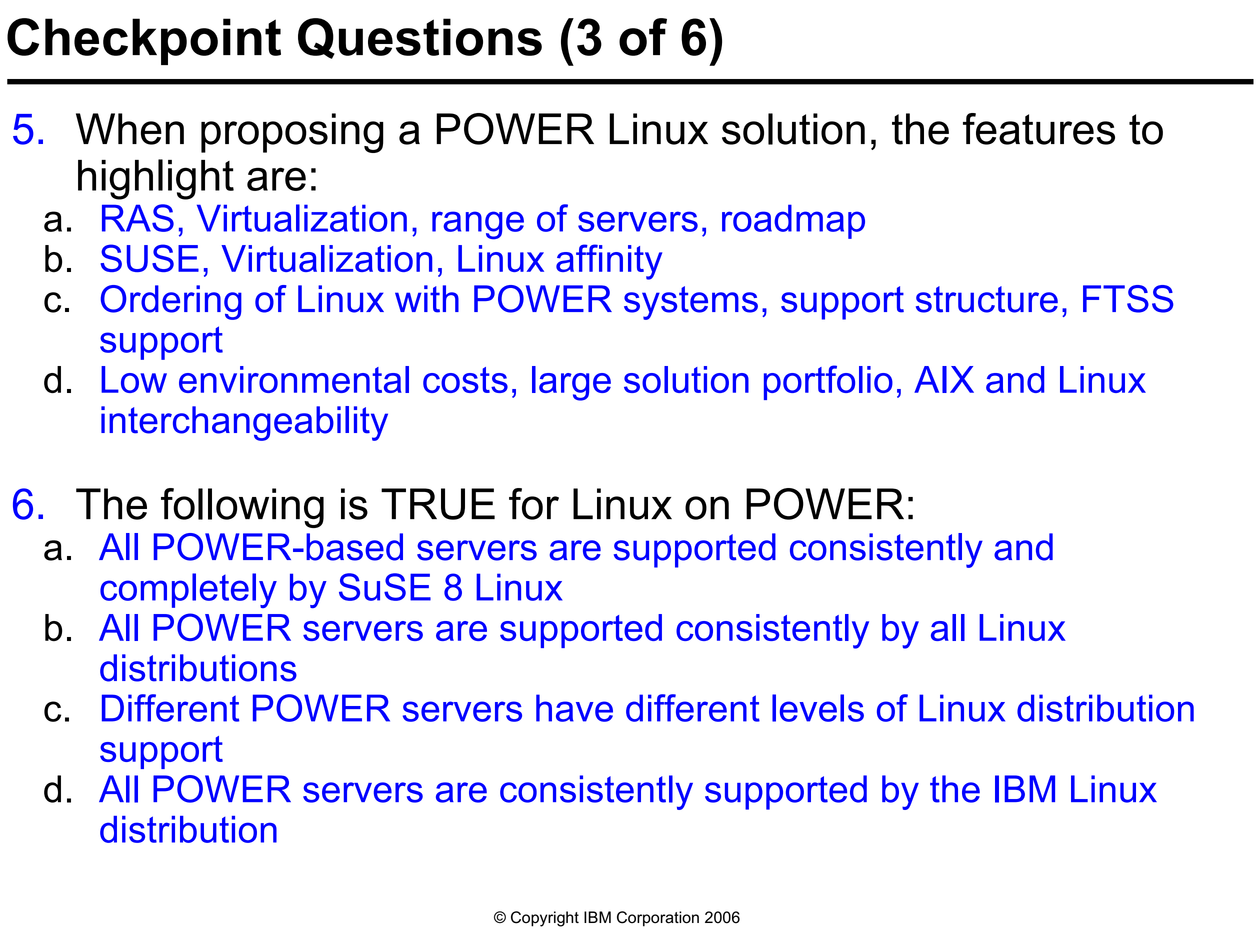

Checkpoint Questions (3 of 6)

5. When proposing a POWER Linux solution, the features to highlight are:
   a. RAS, Virtualization, range of servers, roadmap
   b. SUSE, Virtualization, Linux affinity
   c. Ordering of Linux with POWER systems, support structure, FTSS support
   d. Low environmental costs, large solution portfolio, AIX and Linux interchangeability

6. The following is TRUE for Linux on POWER:
   a. All POWER-based servers are supported consistently and completely by SuSE 8 Linux
   b. All POWER servers are supported consistently by all Linux distributions
   c. Different POWER servers have different levels of Linux distribution support
   d. All POWER servers are consistently supported by the IBM Linux distribution

Figure 7-22. Checkpoint Questions (3 of 6) AM711.0

Notes:

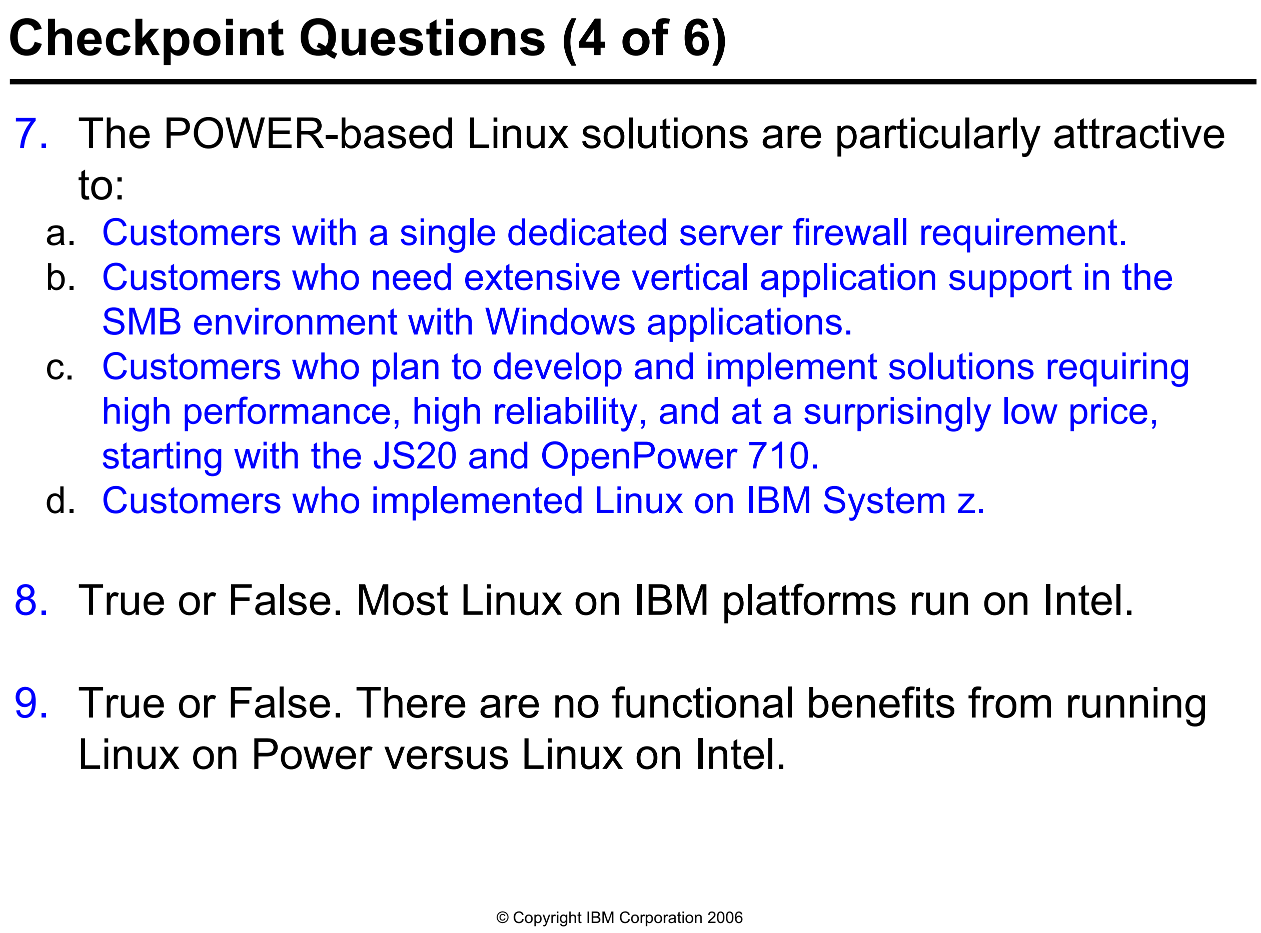

Checkpoint Questions (4 of 6)

7. The POWER-based Linux solutions are particularly attractive to:
   a. Customers with a single dedicated server firewall requirement.
   b. Customers who need extensive vertical application support in the SMB environment with Windows applications.
   c. Customers who plan to develop and implement solutions requiring high performance, high reliability, and at a surprisingly low price, starting with the JS20 and OpenPower 710.
   d. Customers who implemented Linux on IBM System z.

8. True or False. Most Linux on IBM platforms run on Intel.

9. True or False. There are no functional benefits from running Linux on Power versus Linux on Intel.

Figure 7-23. Checkpoint Questions (4 of 6) AM711.0

Notes:

Topic Summary

Having completed this unit, you should be able to:

- Describe the Linux market and position Power architecture systems within that market.
- Describe the market opportunity created by Linux for IBM System p.
- Describe the major features of Linux to clients.
- Describe the Linux on Power benefits versus Linux on Intel.

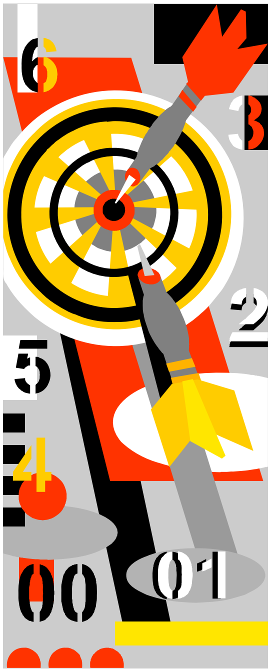

Figure 7-24. Topic Summary AM711.0

Notes:

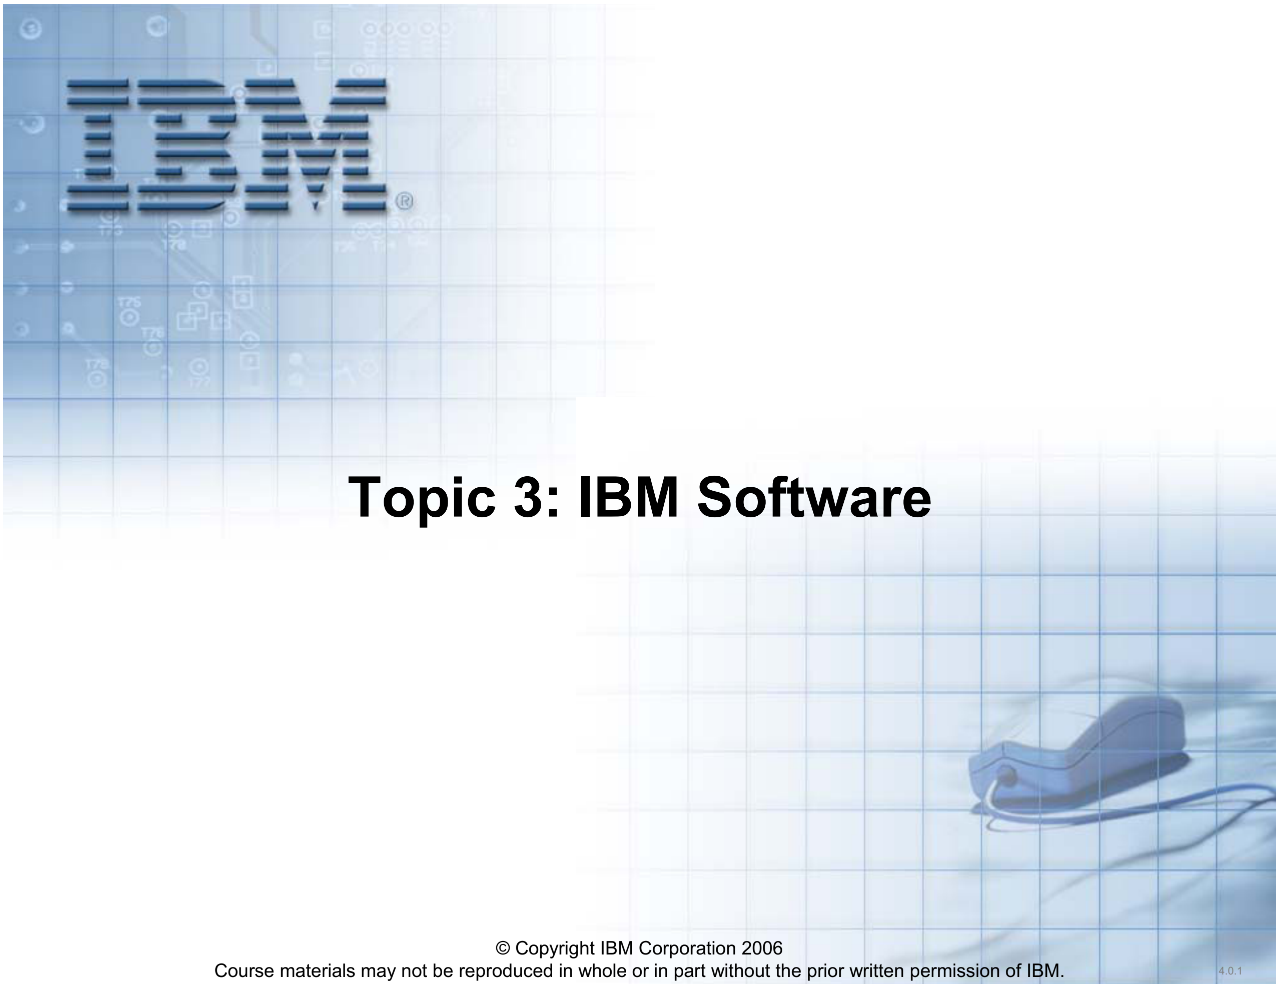

Topic 3: IBM Software

4.0.1

Figure 7-25. Topic 3: IBM Software AM711.0

Notes:

Topic Objectives

- After completing this topic you should be able to:
  - Describe Processor Value Unit (PVU) licensing.
  - Discuss IBM software products.

Figure 7-26. Topic Objectives AM711.0

Notes:

What is Processor-Value Unit Licensing?

- A licensing structure and strategy for per processor priced middleware in Passport Advantage
  - Simple conversion from current licensing
  - Effectively software prices do not change for existing processors

New July 2006

Figure 7-27. What is Processor-Value Unit Licensing? AM711.0

Notes:

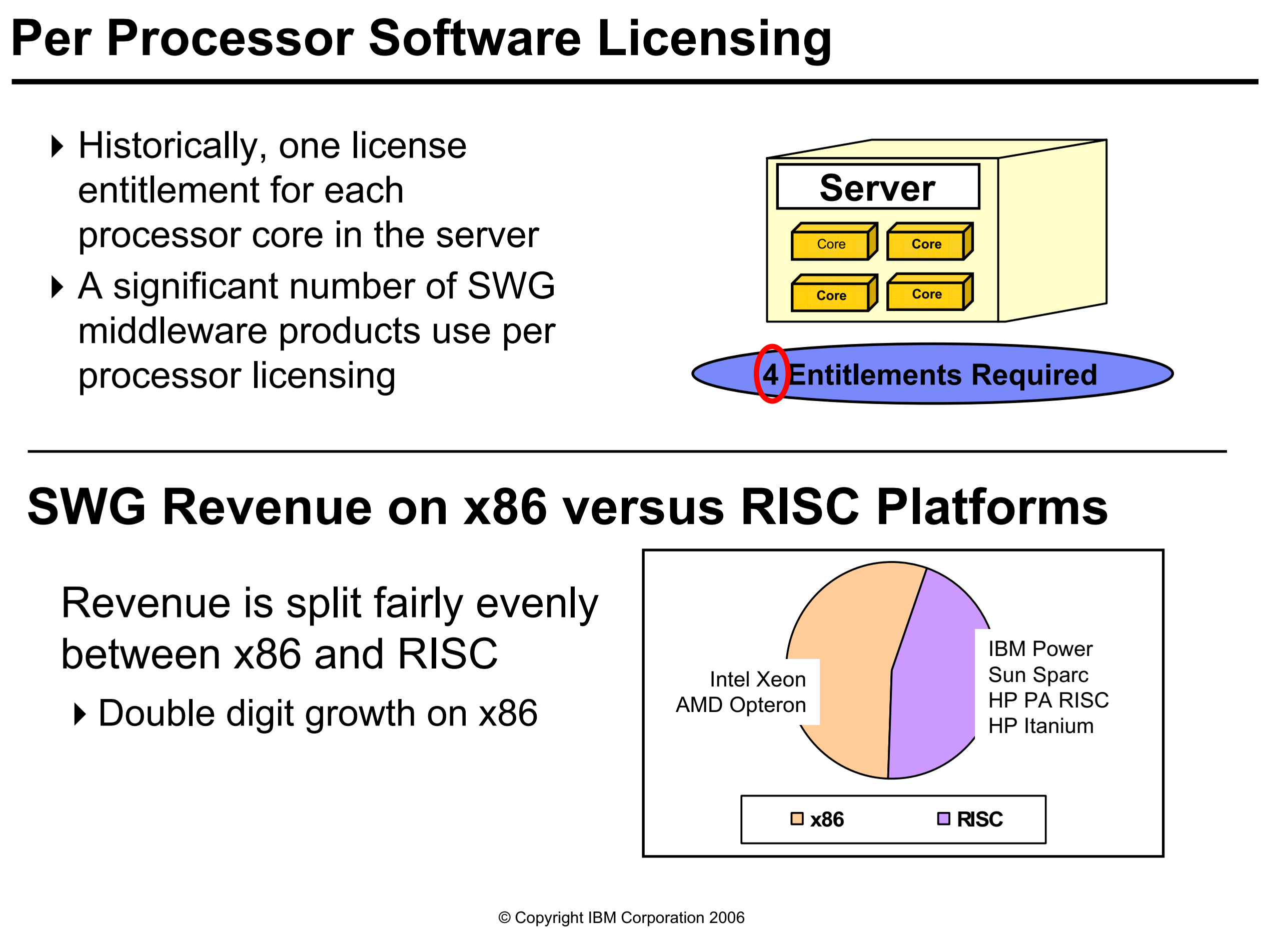

Per Processor Software Licensing

- Historically, one license entitlement for each processor core in the server
- A significant number of SWG middleware products use per processor licensing

Server
Core
Core
Core
Core
4 Entitlements Required

SWG Revenue on x86 versus RISC Platforms

Revenue is split fairly evenly between x86 and RISC

- Double digit growth on x86

Intel Xeon
AMD Opteron
IBM Power
Sun Sparc
HP PA RISC
HP Itanium
x86
RISC

Figure 7-28. Per Processor Software Licensing AM711.0

Notes:

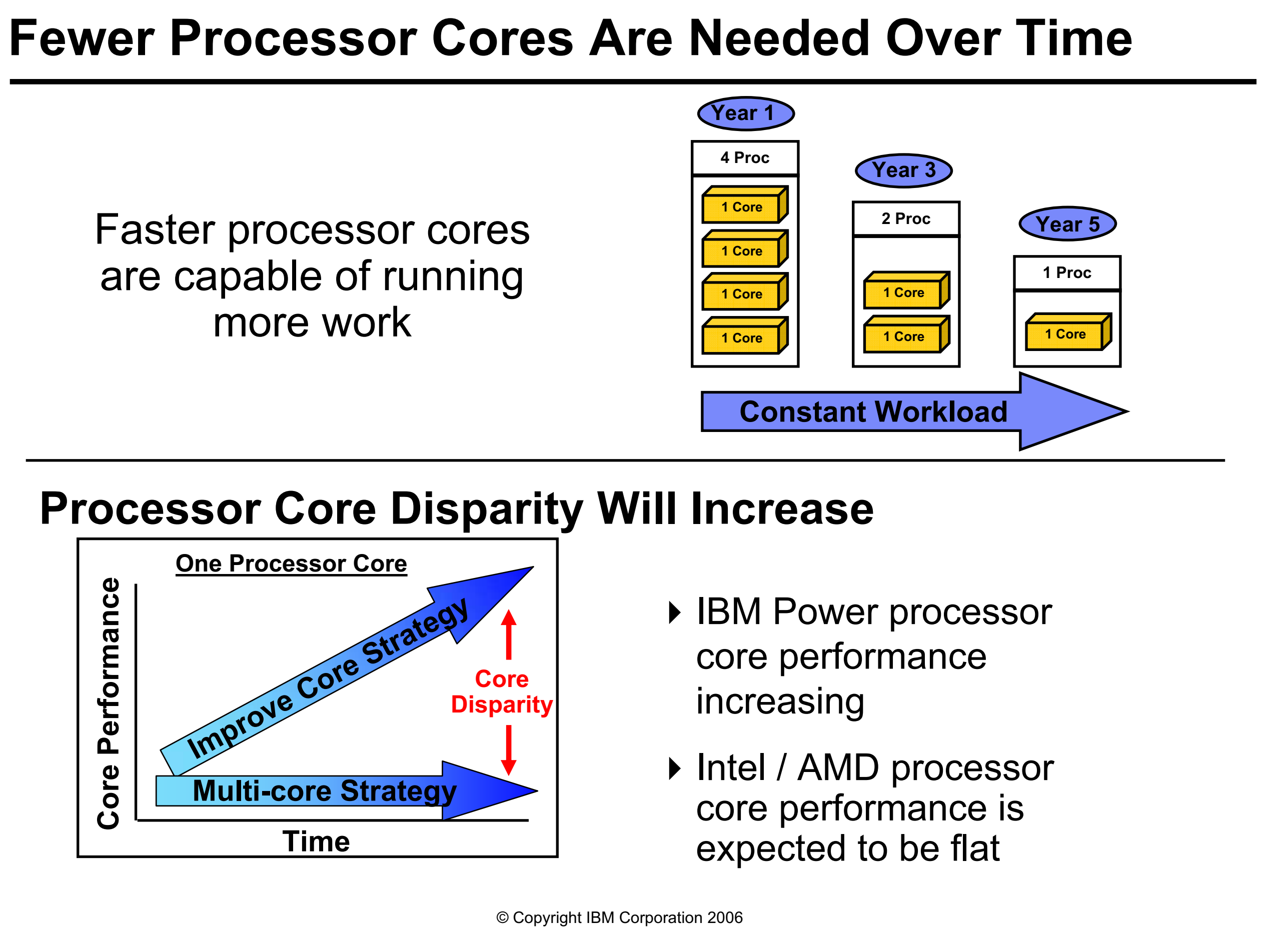

Fewer Processor Cores Are Needed Over Time

Faster processor cores
are capable of running
more work

Processor Core Disparity Will Increase

- IBM Power processor core performance increasing
- Intel / AMD processor core performance is expected to be flat

Figure 7-29. Fewer Processor Cores Are Needed Over Time AM711.0

Notes:

Software Vendors Differentiate by Processor Type

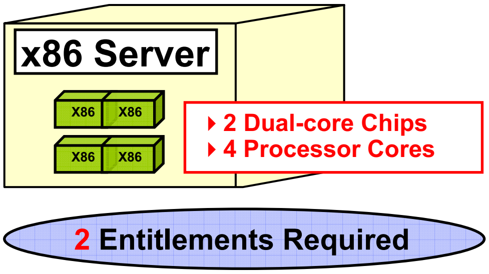

On x86 dual-core chips, 0.50 license entitlements are required per processor

(IBM per processor policy for x86 dual-core is "1 entitlement for 2 cores")

This has added complexity to software licensing.

IBM has three different middleware per processor core license entitlement levels today

	Effective Entitlements Per Processor			
Chip Type	**IBM**	**Microsoft***	**Oracle**	**BEA**
Single-core (All Chips)	1.00	1.00	1.00	1.00
RISC Dual-core	1.00	N/A	0.75	0.50
x86 Dual-core	0.50	0.50	0.50	0.50
RISC Sun T1 Octi-core	0.38	N/A	0.25	0.25

(Result: fractional licenses)

* Microsoft per core entitlements are derived from their per chip licensing.

Figure 7-30. Software Vendors Differentiate by Processor Type AM711.0

Notes:

PVU Provides Flexibility through Granularity

- Value unit licensing at the processor core level
  - Capability to enable sub-capacity licensing
  - Granularity to address changing technology
- Flexibility to create new processor value unit levels
  - Deliver software price performance improvements

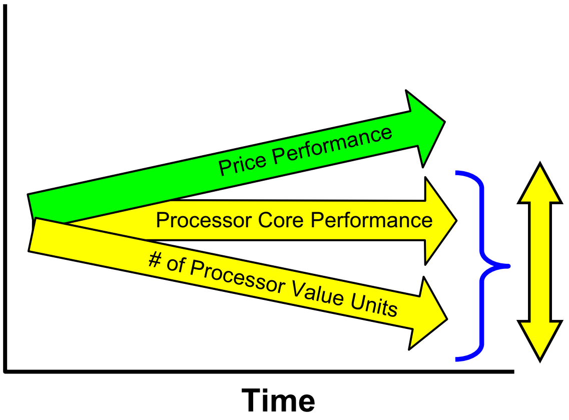

Software price performance improvements will range over time depending upon a number of factors including market conditions

Processor core performance levels will vary between chip vendors
- Processor Value Units relative to core performance

Figure 7-31. PVU Provides Flexibility through Granularity AM711.0

Notes:

Simple Migration of Existing Entitlements

- Existing Maintenance Processor Entitlements Converted to Processor Value Unit Entitlements
  - Current per Processor Entitlements x 100 = Processor Value Units
  - Same conversion factor for all processor types

	"Old" Structure			New Structure
	Effective Entitlements Per Chip	Effective Entitlements Per Processor	Per Processor Migration Conversion Factor	Migrated Value Unit Entitlements Per Processor
Single-core (All Chips)	1.00	1.00	x 100	100
RISC Dual-core	2.00	1.00	x 100	100
x86 Dual-core	1.00	0.50	x 100	50
RISC Sun T1 Octi-core	3.00	0.30 *	x 100	30

* T1 entitlements per processor adjusted with the elimination of fractional licenses

Figure 7-32. Simple Migration of Existing Entitlements AM711.0

Notes:

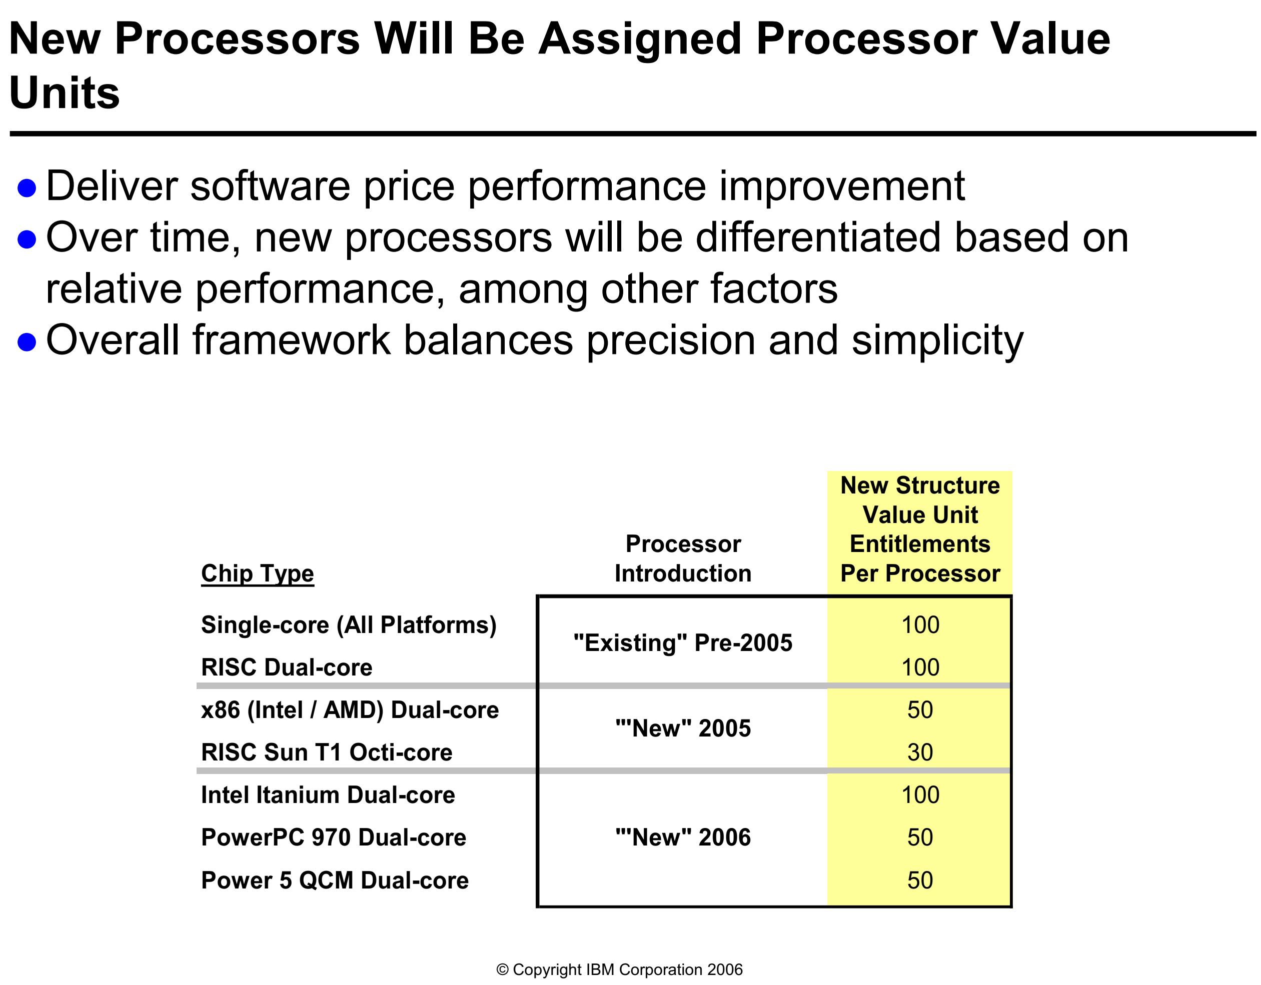

New Processors Will Be Assigned Processor Value Units

- Deliver software price performance improvement
- Over time, new processors will be differentiated based on relative performance, among other factors
- Overall framework balances precision and simplicity

Chip Type	**Processor Introduction**	**New Structure Value Unit Entitlements Per Processor**
Single-core (All Platforms)	**"Existing" Pre-2005**	100
RISC Dual-core		100
x86 (Intel / AMD) Dual-core	**'"New" 2005**	50
RISC Sun T1 Octi-core		30
Intel Itanium Dual-core	**'"New" 2006**	100
PowerPC 970 Dual-core		50
Power 5 QCM Dual-core		50

Figure 7-33. New Processors Will Be Assigned Processor Value Units AM711.0

Notes:

Processor Value Unit Calculator

- Helps determine license quantities
- Passport Advantage Online contains a link to the processor Value Unit Calculator
  - Calculates number of processor value units required for every hardware environment
  - Used to calculate price*

*(Number of value units x price per value unit)

Figure 7-34. Processor Value Unit Calculator AM711.0

Notes:

Topic Summary

- Having completed this topic, you should be able to:
  - Describe Processor Value Unit (PVU) licensing.
  - Discuss IBM software products.

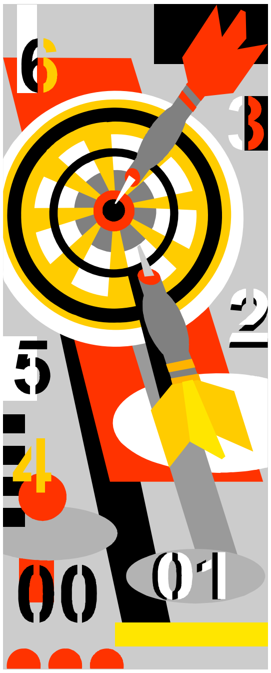

Figure 7-35. Topic Summary AM711.0

Notes:

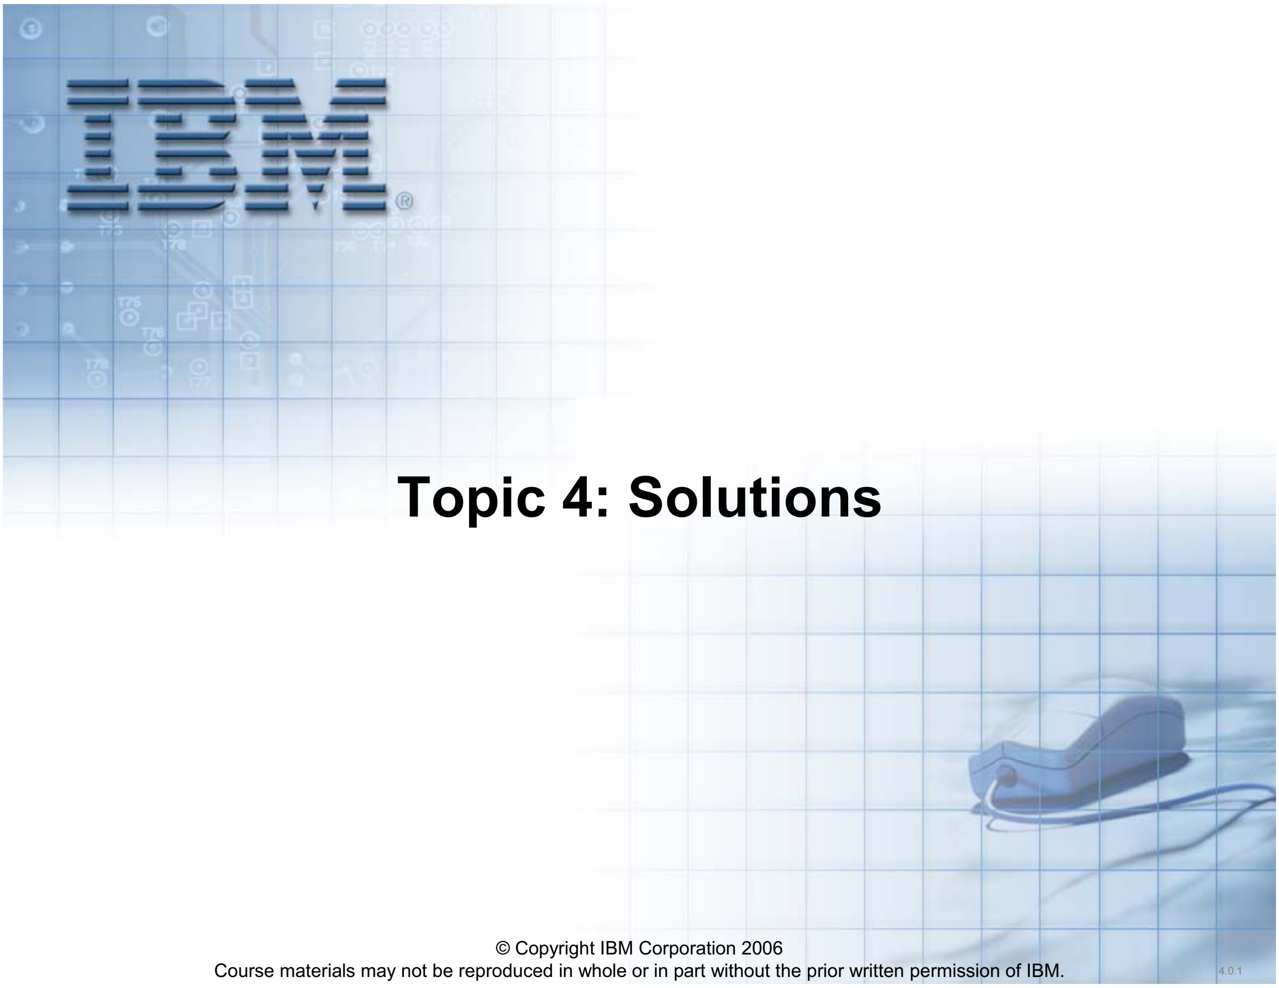

Topic 4: Solutions

Figure 7-36. Topic 4: Solutions AM711.0

Notes:

Topic Objectives

- After completing this unit, you should be able to:
  - List the components of solutions
  - Identify the benefits of solutions for client and seller.
  - Describe the solutions available with System p.
  - Recognize solutions opportunities.
  - Describe the role of the IBM Server Consolidation Factory in solution selling.
  - Describe the benefits to client and seller of using IBM Global Financing (IGF) as part of solution selling.

Figure 7-37. Topic Objectives AM711.0

Notes:

What Are Solutions?

Linux - or- **AIX 5L**

+

System p Servers

+

System Storage (selected solutions)

+

DB2, WebSphere (selected solutions)

+

IBM Services and Financing (selected solutions)

+

High Value ISV Applications

=

Solutions

ISV Solutions
- **On our Web site**
- **Solution Brief exists**

*Solution Offerings**
- **Tested and sized by IBM**
- **At least three (SML) recommended configurations**
- **Sales Kit available**

*Solution Editions**
- **Tested and sized by IBM**
- **At least three (SML) specially priced recommended configs.**
- **Sales Kit available**

Featured Solution
- **Any of the above featured on our Web site**

* Focus of high-touch G-T-M Efforts

Figure 7-38. What Are Solutions? AM711.0

Notes:

Why Create Solution Editions and Solution Offerings?

For our clients:

- Pre-configuration speeds time to implementation which **improves** project **ROI.**
- Integration and testing **reduces** implementation **risk.**
- IBM doing the work **relieves** the level and quantity of **personnel skill.**
- Financing for the entire solution, including IBM and non-IBM hardware, software and services, all under a single contract **simplifies** the **purchasing** process.

Figure 7-39. Why Create Solution Editions and Solution Offerings? AM711.0

Notes:

Why Create Solution Editions and Solution Offerings?

For our sellers:

- The customer value makes it **easier to sell.**
- Clients buy servers to do something. Selling that something as a solution puts you **closer to** their **buying behavior.**
- You can **add value** around the solutions – services, and so on – **with confidence** that we will provide programs, collateral and enablement that make it easier to get ROI on your investment.

Figure 7-40. Why Create Solution Editions and Solution Offerings? AM711.0

Notes:

IBM System p5 Solution Offerings

- 10 new solutions
  - Including Solution Editions for SAP BI, Oracle E-BS and JDE E1
- Pre-sized configurations
  - Small enough and large enough for almost every client
- Validated by experts
  - IBM and the ISV stand behind the solution
- Easy for our clients to implement
  - We have done a lot of the technical work for them

NEW!
July 2006

Figure 7-41. IBM System p5 Solution Offerings AM711.0

Notes:

Introducing Solutions from System p

Integrated and tested to speed time to implementation for improved ROI

IBM System p5 Solution Edition for SAP NetWeaver Business Intelligence (BI)	AIX L	redhat	suse
IBM System p5 Solution Edition for Oracle E-Business Suite	AIX L		
IBM System p5 Solution Edition for Oracle's JD Edwards EnterpriseOne	AIX L		
SAS Forecast Server for AIX 5L on IBM POWER processor-based systems	AIX L		
IBM and Misys Express Solution for Healthcare Providers	AIX L		
IBM DB2 9 for System p5	AIX L	redhat	suse
Sybase Mobility Solution for IBM System p5		redhat	suse
Oracle Real Application Clusters for IBM System p5	AIX L	redhat	suse
IBM System p5 for SAP's Enterprise SOA	AIX L	redhat	suse
IBM BladeCenter JS21 for Web serving farms		redhat	suse
Rational PurifyPlus for AIX 5L on IBM POWER processor-based systems	AIX L		

Figure 7-42. Introducing Solutions from System p AM711.0

Notes:

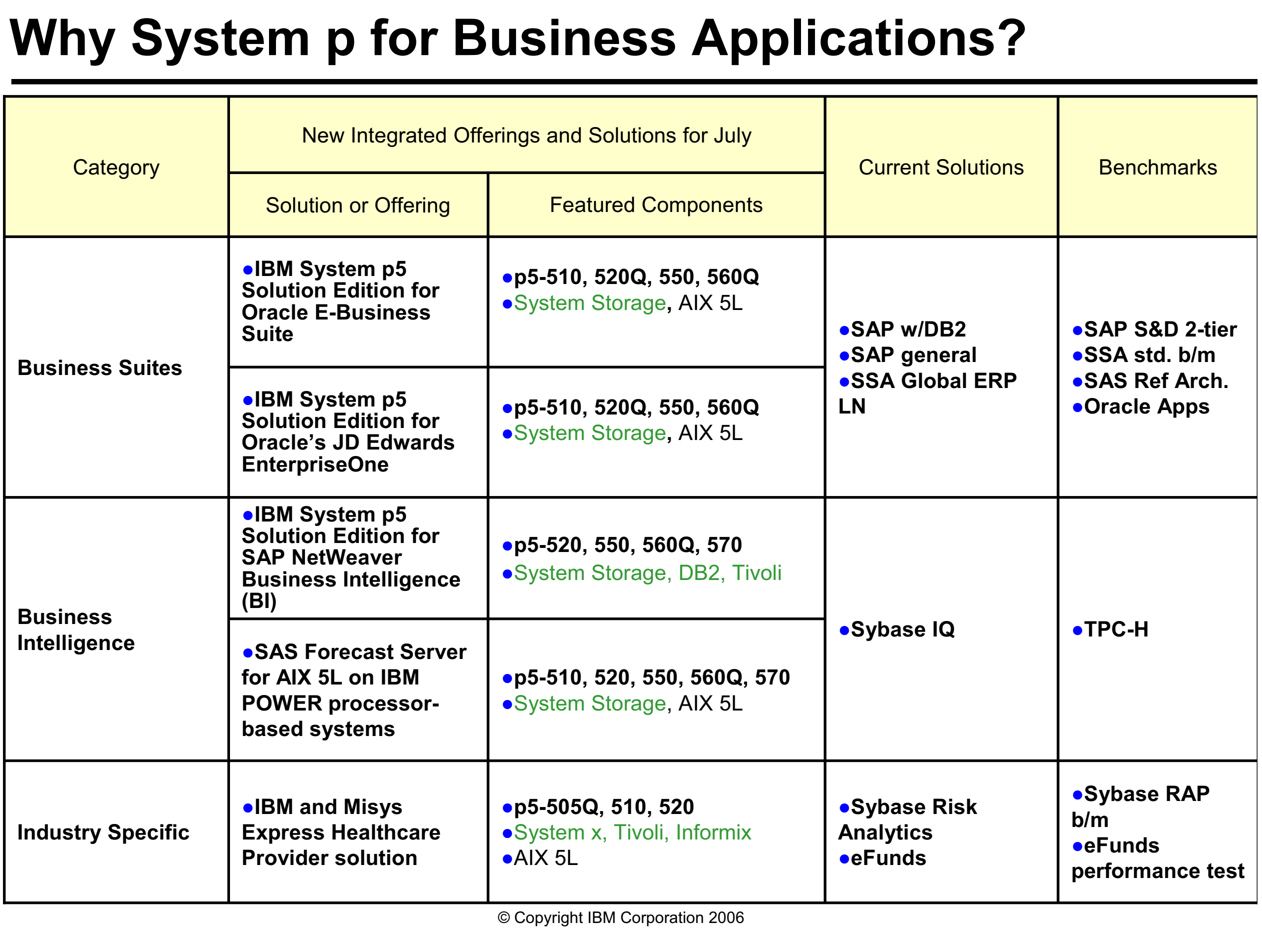

Why System p for Business Applications?

Category	New Integrated Offerings and Solutions for July		Current Solutions	Benchmarks
	Solution or Offering	Featured Components		
Business Suites	**●IBM System p5 Solution Edition for Oracle E-Business Suite**	**●p5-510, 520Q, 550, 560Q** ●System Storage, AIX 5L	**●SAP w/DB2** **●SAP general** **●SSA Global ERP LN**	**●SAP S&D 2-tier** **●SSA std. b/m** **●SAS Ref Arch.** **●Oracle Apps**
	●IBM System p5 Solution Edition for Oracle's JD Edwards EnterpriseOne	**●p5-510, 520Q, 550, 560Q** ●System Storage, AIX 5L		
Business Intelligence	**●IBM System p5 Solution Edition for SAP NetWeaver Business Intelligence (BI)**	**●p5-520, 550, 560Q, 570** ●System Storage, DB2, Tivoli	**●Sybase IQ**	**●TPC-H**
	●SAS Forecast Server for AIX 5L on IBM POWER processor-based systems	**●p5-510, 520, 550, 560Q, 570** ●System Storage, AIX 5L		
Industry Specific	**●IBM and Misys Express Healthcare Provider solution**	**●p5-505Q, 510, 520** ●System x, Tivoli, Informix ●AIX 5L	**●Sybase Risk Analytics** **●eFunds**	**●Sybase RAP b/m** **●eFunds performance test**

Figure 7-43. Why System p for Business Applications? AM711.0

Notes:

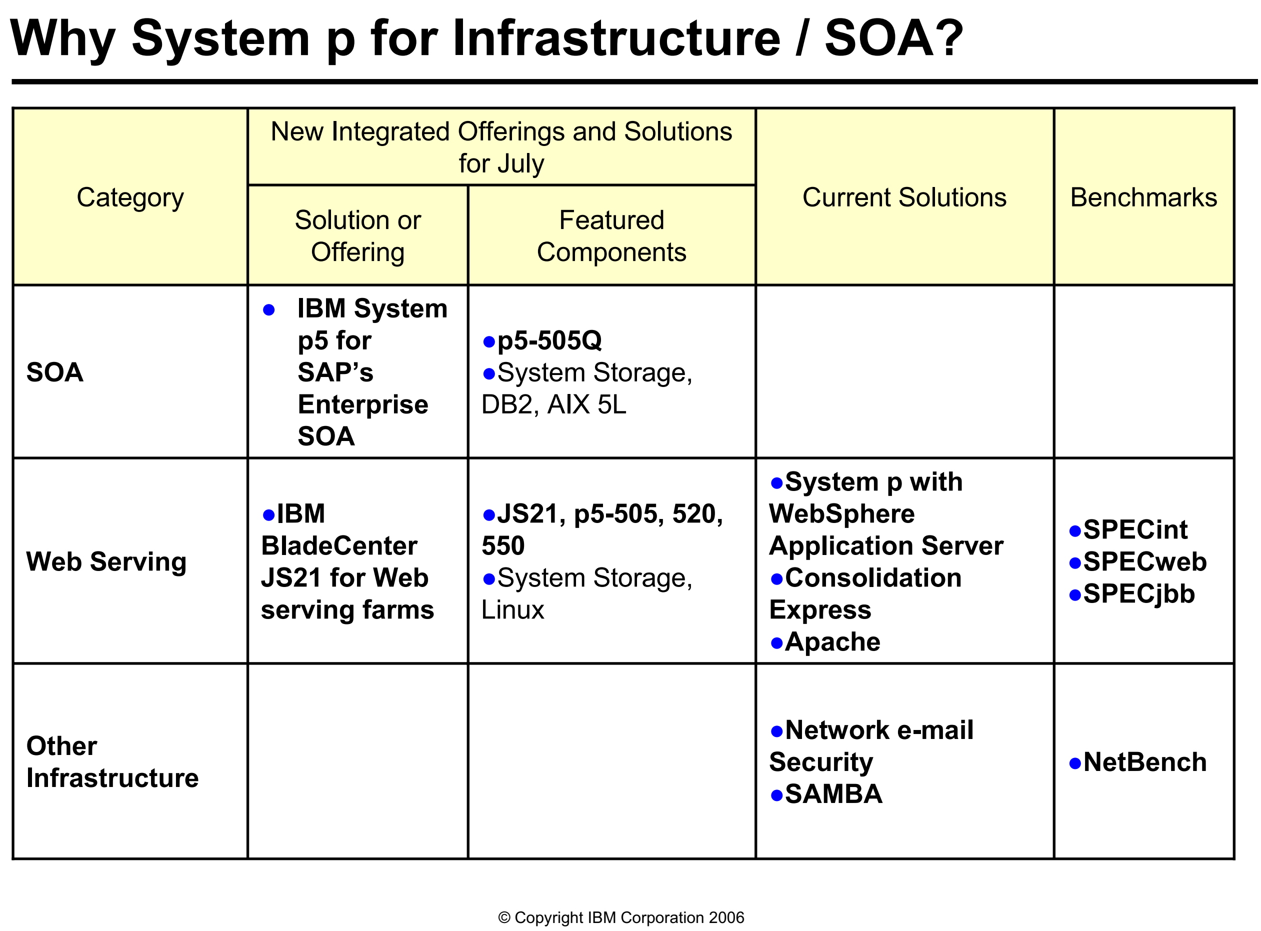

Why System p for Infrastructure / SOA?

Category	New Integrated Offerings and Solutions for July		Current Solutions	Benchmarks
	Solution or Offering	Featured Components		
SOA	● **IBM System p5 for SAP's Enterprise SOA**	●**p5-505Q** ●System Storage, DB2, AIX 5L		
Web Serving	●**IBM BladeCenter JS21 for Web serving farms**	●**JS21, p5-505, 520, 550** ●System Storage, Linux	●**System p with WebSphere Application Server** ●**Consolidation Express** ●**Apache**	●**SPECint** ●**SPECweb** ●**SPECjbb**
Other Infrastructure			●**Network e-mail Security** ●**SAMBA**	●**NetBench**

Figure 7-44. Why System p for Infrastructure / SOA? AM711.0

Notes:

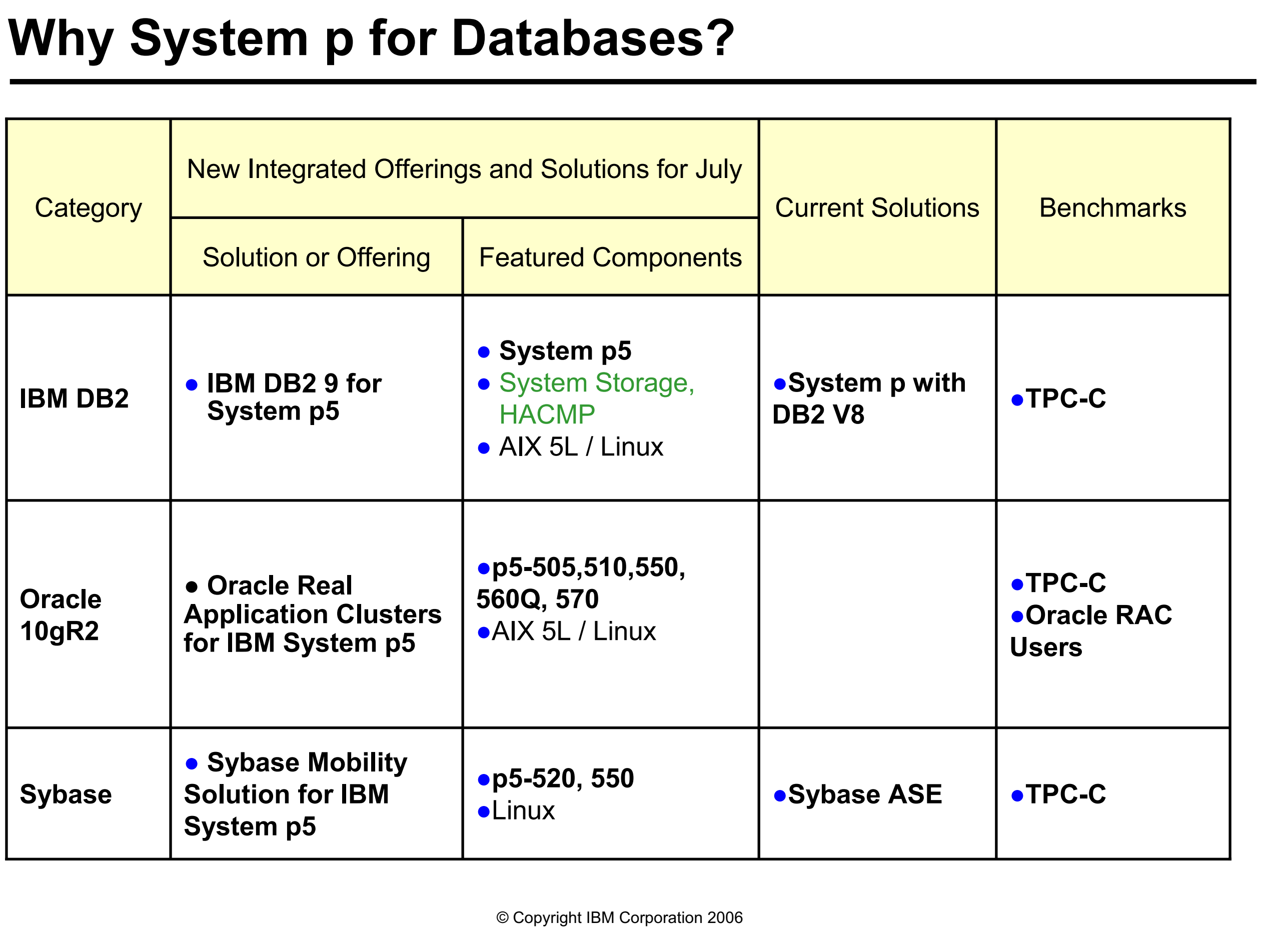

Why System p for Databases?

Category	New Integrated Offerings and Solutions for July		Current Solutions	Benchmarks
	Solution or Offering	Featured Components		
IBM DB2	● **IBM DB2 9 for System p5**	● **System p5** ● System Storage, HACMP ● AIX 5L / Linux	●**System p with DB2 V8**	●**TPC-C**
Oracle 10gR2	● **Oracle Real Application Clusters for IBM System p5**	●**p5-505,510,550, 560Q, 570** ●AIX 5L / Linux		●**TPC-C** ●**Oracle RAC Users**
Sybase	● **Sybase Mobility Solution for IBM System p5**	●**p5-520, 550** ●Linux	●**Sybase ASE**	●**TPC-C**

Figure 7-45. Why System p for Databases? AM711.0

Notes:

Why System p for HPC?

Category	New Integrated Offerings and Solutions for July		Current Solutions	Benchmarks
	Solution or Offering	Featured Components		
Life Sciences			**●Bioinformatics** **●Agilent / Genomics** **●Termo Electron/Proteomics** **●Waters/Proteomics** **●CPMD and GAMESS**	**HMMER, Blast, Fasta CPMD, GAMES LINPACK**
Petro	**●IBM BladeCenter PowerPC Solution for Seismic Data Processing and Imaging**	**●JS21 Blade** ●Linux		**CGG**
CAE	**●IBM BladeCenter JS21 Solution for Computer Aided Engineering (CAE)**	**●JS21 Blade** ●AIX 5L, Linux		
Grid and Clusters	**●HPC Cluster-in-a-BladeCenter (3Q)**	**●JS21 Blade** ●System Storage ●Linux	**●Grid-and-grow Express**	

Figure 7-46. Why System p for HPC? AM711.0

Notes:

IBM Server Consolidation Factory for System p

- Product focus:
  - IBM System p5 590 and 595 servers with Advanced POWER Virtualization
  - IBM System Storage DS8000
  - High Availability Cluster Multiprocessing (HACMP)
  - Tivoli Usage and Accounting Manager
- Service options:
  - Initial consulting to determine the project scope, potential TCO and business benefits
  - Detailed planning and system design
  - Implementation services
- Financing options:
  - Trade-in program to help reduce up-front costs of consolidation
  - Lease and purchase options to help meet TCO objectives, conserve cash, enable utility on-off server usage

Figure 7-47. IBM Server Consolidation Factory for System p AM711.0

Notes:

Opportunities for Server Consolidation

- Key industries:
  - Government, Banking, Healthcare, Computer Services, Travel and Transportation
- Customer size: Large enterprise
  - Clients who are considering a major data center refresh or upgrade
    - Major software change or application expansion
    - Reducing number of server vendors' platforms
    - Data center is moving, consolidating (acquisition) or being replicated (backup or divestiture)
  - Clients who want to improve operational efficiency and reduce costs without sacrificing service levels on mission-critical applications such as ERP, CRM or large database applications
- **Decision-makers: CEO, CIO, or CFO**
- **Influencers: CIO, IT VP, or IT Director**

Figure 7-48. Opportunities for Server Consolidation AM711.0

Notes:

IBM Server Consolidation Factory for System p

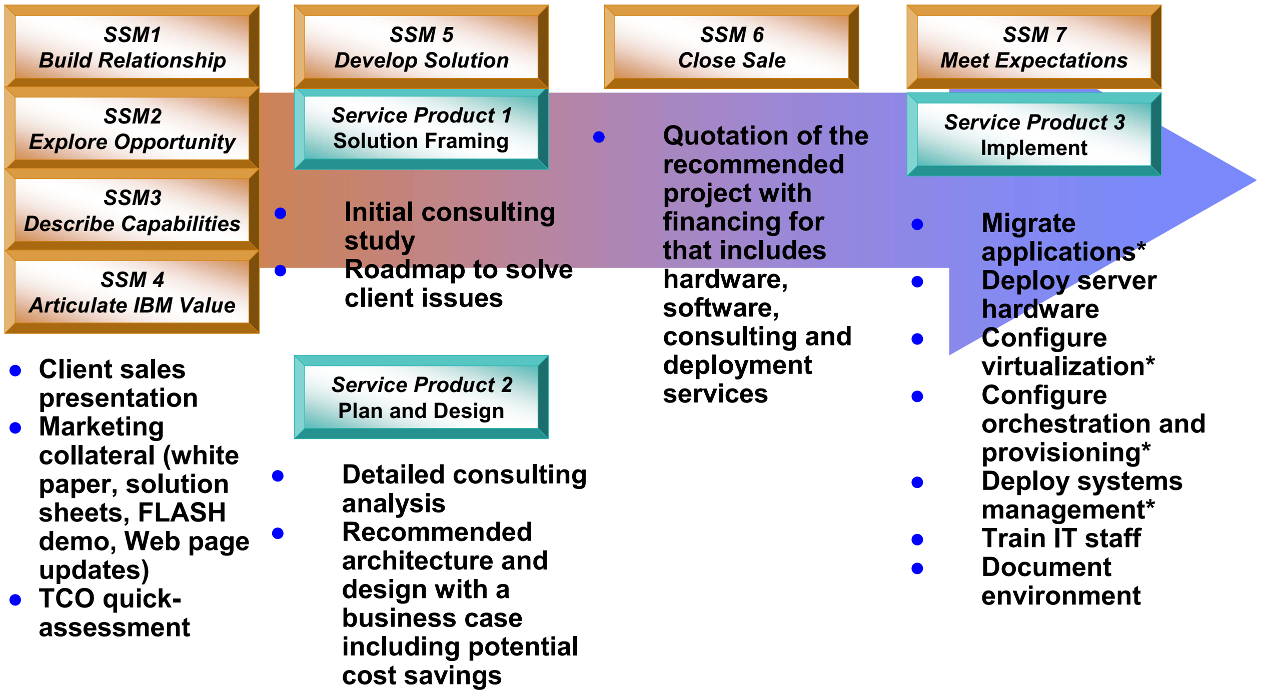

SSM1
Build Relationship

SSM2
Explore Opportunity

SSM3
Describe Capabilities

SSM 4
Articulate IBM Value

- Client sales presentation
- Marketing collateral (white paper, solution sheets, FLASH demo, Web page updates)
- TCO quick-assessment

SSM 5
Develop Solution

Service Product 1
Solution Framing

- Initial consulting study
- Roadmap to solve client issues

Service Product 2
Plan and Design

- Detailed consulting analysis
- Recommended architecture and design with a business case including potential cost savings

SSM 6
Close Sale

- Quotation of the recommended project with financing for that includes hardware, software, consulting and deployment services

SSM 7
Meet Expectations

Service Product 3
Implement

- Migrate applications*
- Deploy server hardware
- Configure virtualization*
- Configure orchestration and provisioning*
- Deploy systems management*
- Train IT staff
- Document environment

*If applicable

Figure 7-49. IBM Server Consolidation Factory for System p AM711.0

Notes:

Consulting and Deployment Services

- Consulting with IBM Global Technology Services helps clients understand and manage costs
  - **Project design report** including business goals, recommended targets for hardware consolidation and application migration, projected Total Cost of Ownership, and roadmap for business transition
  - **Plan and design report** including detailed analysis of current environment, design for hardware consolidation and application migration, and project plan for rollout
- Implementation services help clients reduce risk to business operations
  - **Application migration and hardware deployment** including partitioning and virtualization configuration, and verification testing
  - **Systems management tool deployment**
  - **IT staff training** including onsite training on hardware and systems management tools, and vouchers for partitioning and virtualization classes

Figure 7-50. Consulting and Deployment Services AM711.0

Notes:

IBM Tivoli Usage and Accounting Manager

- **Helps clients maintain business units' satisfaction with services when workloads are consolidated**
  - Allocate/distribute or charge IT costs to the users, cost centers, applications
  - Align business operations and IT costs
- **Works in consolidated environments with physical and virtualized resources**
  - Supports multiple platforms including Mainframe, UNIX, AIX 5L, Linux and Windows
  - Supports multiple sub-systems and workloads including Database, Web, Storage and Networks
- **Designed to be an integral part of your client's financial reporting systems**

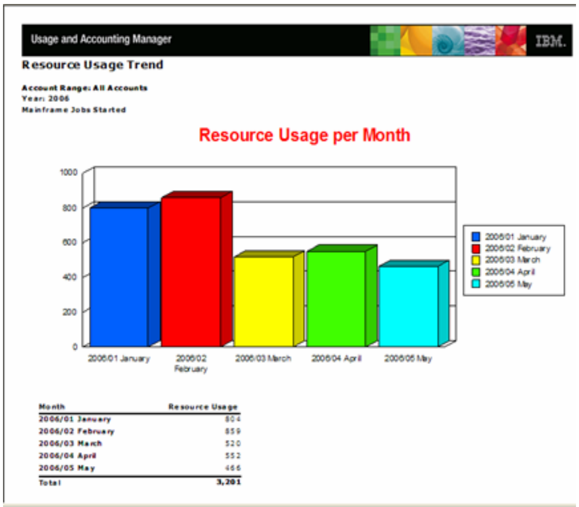

Resource Usage Trend

Account Range: All Accounts
Year: 2006
Mainframe Jobs Started

Month	Resource Usage
2006/01 January	804
2006/02 February	859
2006/03 March	520
2006/04 April	552
2006/05 May	466
Total	3,201

Unique accounting features in AIX 5L V5.3 allow System p clients to benefit from detailed reporting not available to other UNIX platforms.

Figure 7-51. IBM Tivoli Usage and Accounting Manager AM711.0

Notes:

This demonstrates the various options of LoP and shows that POWER continues to have higher support when Lintel stops, clustering the only method for higher capacity.

Funding Solution Purchases

- **All investments need financing, either with internal funds or external sources**
- **Financing helps clients match their solution cost outlays to expected ROI benefits and/or budget requirements**
- **IGF can finance large- and small-scale solution purchases**
  - Includes IBM and non-IBM hardware, software and services
- **Highly-customized financing/lease structures and payment terms tailored to client needs**
  - Flexible repayment options for both upfront and ongoing operating costs
  - Special low-rate financing, deferred payment options
  - IT spending aligned to business requirements

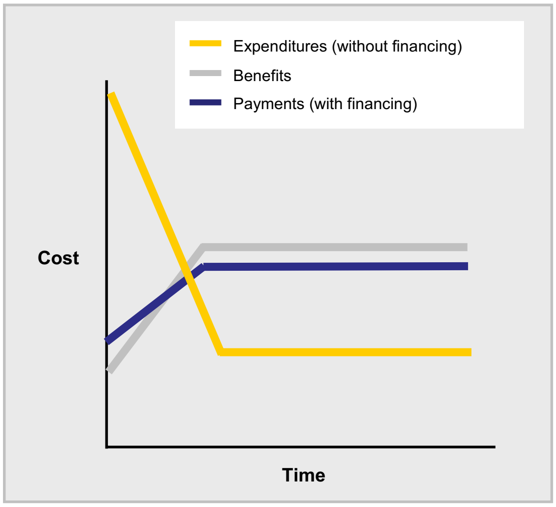

Figure 7-52. Funding Solution Purchases AM711.0

Notes:

The Client Benefits of IT Financing

Financing can be smarter than outright purchase

- **Match cash flow outlays with expected benefits**
- **Lowers total cost of ownership (TCO)**
- **Increases flexibility in the marketplace**
- **Transfers the risk of obsolescence to an external organization**

Financing can be simpler than outright purchase

- **Gives clients the ability to finance a multifaceted IT solution through a single contract**

Financing offers the opportunity to partner with an expert in the field

- **Expert advice and assistance from a partner dedicated to your success**
- **A financial plan for the life of your solution**
- **Help with evaluating lease versus purchase decisions**

Figure 7-53. The Client Benefits of IT Financing AM711.0

Notes:

The Benefits of Choosing IBM Global Financing

Smarter financial decisions

IBM Global Financing can help conserve cash, lower total cost of ownership and support your overall financial objectives

Customized to your needs

IBM Global Financing provides financing solutions tailored to your unique business needs

Simpler experience

IBM Global Financing can help you simplify your entire financing process

Partnership across the lifecycle

IBM Global Financing is a proven partner dedicated to your success across the lifecycle of your information technology investments

One-stop financing

IBM Global Financing offers complete financing services that help optimize value from your investment in IBM and non-IBM IT

Figure 7-54. The Benefits of Choosing IBM Global Financing AM711.0

Notes:

Financing Solutions – ISV and IBM Seller Perspectives

- For ISVs
  - Close deals that may never otherwise have happened due to budget challenges/deferrals
  - Improve ISV’s DSO by utilizing IGF – improve ISV cash position
  - Increase follow on business through availability of additional funds
  - Bottom line/bundled solution facilitates control over pricing metrics such that an ISV can maximize profits while more effectively responding to competitive pressures
- For IBM sellers
  - Sales uplift of 12% to 25% on IGF financed deals
  - Ongoing reasons of call
    - Mid-lease upgrades
      - Often, clients can upgrade with little or no change to their monthly payment
    - Client must make decision at end of lease
  - Keeps out competition – IGF owns the asset

Figure 7-55. Financing Solutions – ISV and IBM Seller Perspectives AM711.0

Notes:

Checkpoint Questions (5 of 6)

10. Solution offerings add value because they are tested and sized by IBM and include at least three (SML) recommended configurations.
    a. True
    b. False

11. Reduced complexity, lower price/performance and the ability to run more workloads on fewer servers are benefits of server consolidation.
    a. True
    b. False

12. The IBM System p solution offering consists entirely of IBM products and services.
    a. True
    b. False

Figure 7-56. Checkpoint Questions (5 of 6) AM711.0

Notes:

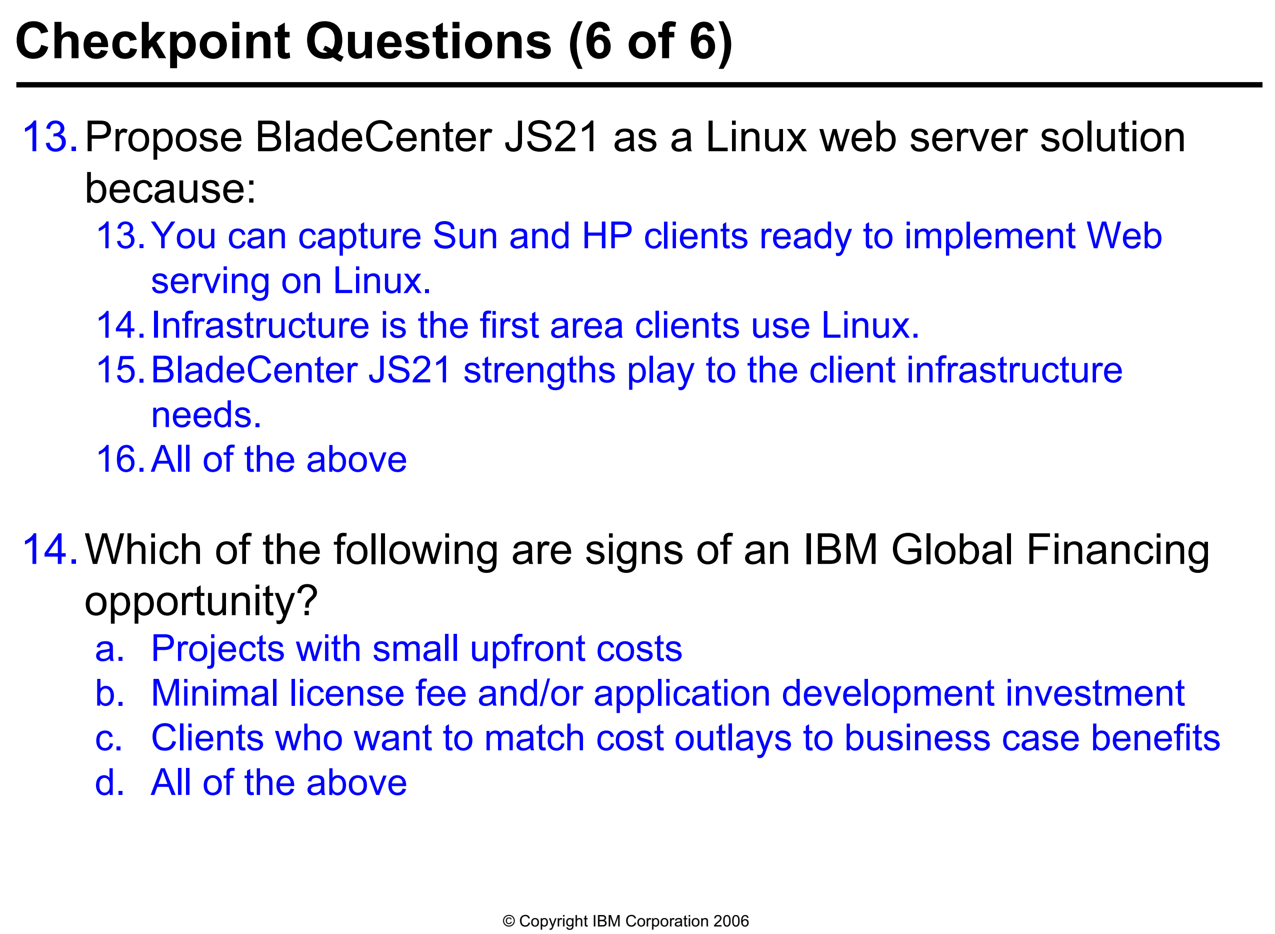

Checkpoint Questions (6 of 6)

13. Propose BladeCenter JS21 as a Linux web server solution because:
    13. You can capture Sun and HP clients ready to implement Web serving on Linux.
    14. Infrastructure is the first area clients use Linux.
    15. BladeCenter JS21 strengths play to the client infrastructure needs.
    16. All of the above

14. Which of the following are signs of an IBM Global Financing opportunity?
    a. Projects with small upfront costs
    b. Minimal license fee and/or application development investment
    c. Clients who want to match cost outlays to business case benefits
    d. All of the above

Figure 7-57. Checkpoint Questions (6 of 6) AM711.0

Notes:

Topic Summary

- Having completed this topic, you should be able to:
  - List the components of solutions.
  - Identify the benefits of solutions for client and seller.
  - Describe the solutions available with System p.
  - Recognize solutions opportunities.
  - Describe the role of the IBM Server Consolidation Factory in solution selling.
  - Describe the benefits to client and seller of using IBM Global Financing (IGF) as part of solution selling.

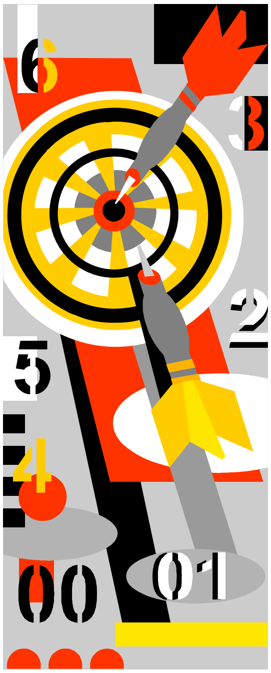

Figure 7-58. Topic Summary AM711.0

Notes:

Unit Summary

Having completed this unit, you should be able to:

- Discuss the benefits of AIX.
- Describe the benefits of Linux on Power.
- Discuss IBM software products and solutions.

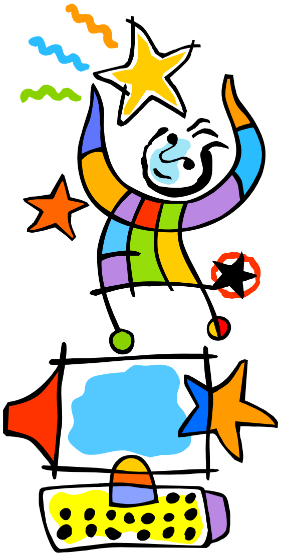

Figure 7-59. Unit Summary AM711.0

Notes:

Unit 8. Competition

What This Unit Is About

This unit describes the competitive advantages of System p.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Explain how our features and capabilities provide competitive advantages
- Position the System p products against the competition

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

After completing this unit, you should be able to:

- Explain how our features and capabilities provide competitive advantages.
- Position the System p products against the competition.

Figure 8-60. Unit Objectives AM715.0

Notes:

Competitive Features of IBM System p

- Performance
- RAS Features
- Virtualization
- Price/performance

Figure 8-61. Competitive Features of IBM System p AM715.0

Notes:

IBM Continues to Lead in UNIX Servers

- But the market complexion is changing
  - Target server market, UNIX + Linux, is growing 2.6% thru 2009
    - UNIX is flat/declining
    - Linux growing 13% year to year.
  - Server price bands are cascading down
    - Performance enabling mid-range SMPs to substitute high-end
    - Substitution and scale-out enabling Linux and UNIX low-end growth

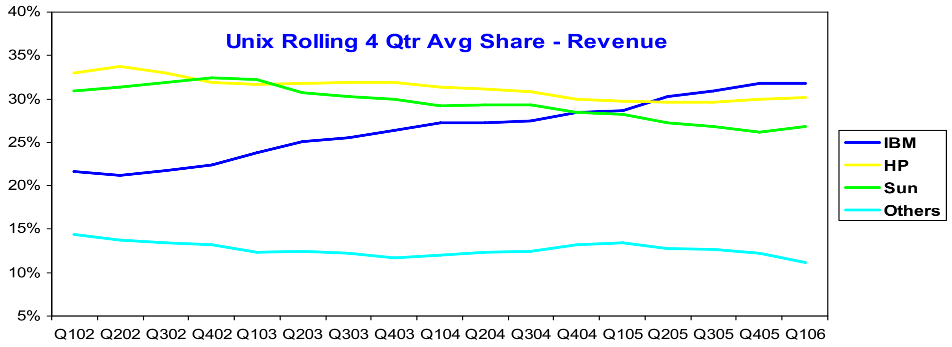

Figure 8-62. IBM Continues to Lead in UNIX Servers AM715.0

Notes:

Our Competitors Are Responding with New Offerings

- HP and Sun pursuing two paths:
  - Refreshing traditional platforms (UltraSPARC IV+, Niagara and HP Itanium2 Montecito)
    - Introducing multi-core systems with increased aggregate performance; will achieve quad-core by 1H07
  - Pushing Linux and Opteron as substitute for UNIX/RISC at the low end
    - Sun Opteron: 1Q: +97% YTY, 9% of total server revenue with 74% running Linux (27% of units)
    - HP Opteron: 1Q: +247% YTY, 12% of total server sales (59% Windows® / 35% Linux)
- Sun gaining new clients via try and buys
  - Niagara - 60% shipping out to new clients, with 50% conversion rate
  - Extended to Opteron line
- Both HP and Sun reinvigorating channels

Figure 8-63. Our Competitors Are Responding with New Offerings AM715.0

Notes:

Sun Opteron Product Roadmap for 2006

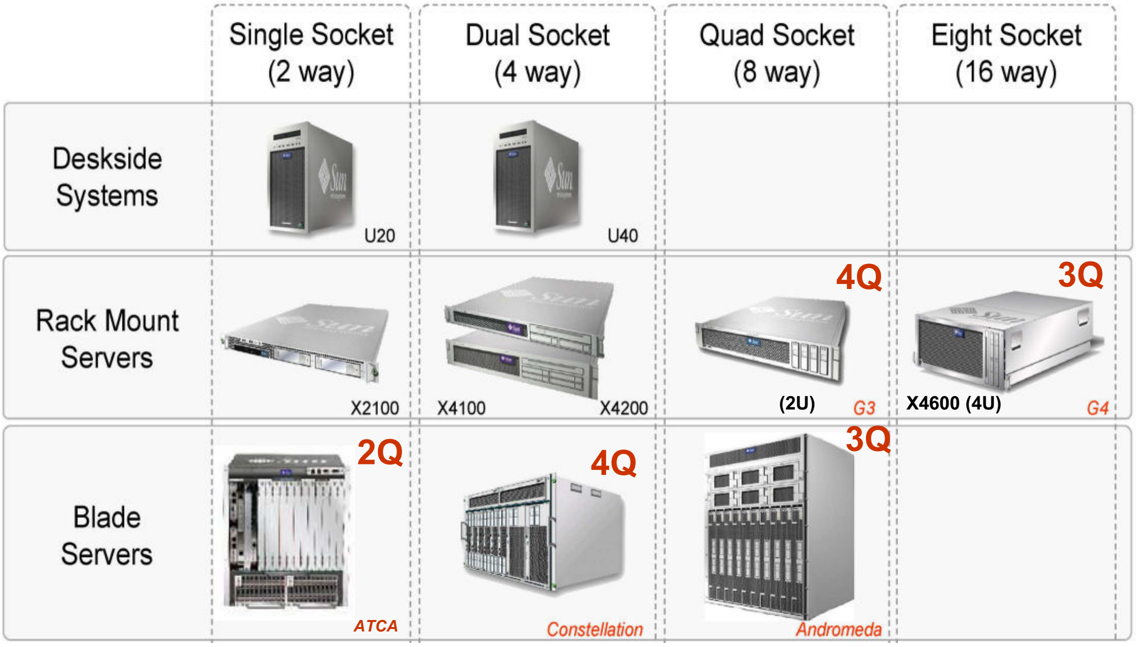

	Single Socket (2 way)	Dual Socket (4 way)	Quad Socket (8 way)	Eight Socket (16 way)
Deskside Systems	U20	U40		
Rack Mount Servers	X2100	X4100 X4200	4Q (2U) G3	3Q X4600 (4U) G4
Blade Servers	2Q ATCA	4Q Constellation	3Q Andromeda	

Figure 8-64. Sun Opteron Product Roadmap for 2006 AM715.0

Notes:

Sun X4600 Response – Attack

- What did Sun announce?
  - First modular 8-socket 4u x64 server with 2x the performance of rival systems
  - Record benchmarks for HPC workloads
  - Upgrade to future quad core processors with "uniboard" design
  - Positioned and compared to other x86 offerings: HP DL585 and IBM x3950
  - Consolidation of x86 servers for a wide range of applications
  - Install base targets are Sun Fire v490/v880/v890

- System p Positioning
  - In the Sun 4- to 8-core sweet spot, the IBM System p5 550 and 550Q provide a better commercial application platform
  - Lead with 560Q when competing against the 16-core X4600

Figure 8-65. Sun X4600 Response – Attack AM715.0

Notes:

Sun Blade X8400 versus JS21 Blade with BladeCenter H

JS21: 2X better SWaP at half the price of Sun Blade X8400*

- Sun Blade X8400
  - Dual-core AMD Opteron 2.6 GHz
  - 10 4-socket blades (80 cores) in 19U chassis
    - 6225 watts / 21227 BTU for 10 blades
    - 622 watts / 2122 BTU per blade
  - Chassis: six 3000W power supplies (N+N)
  - SWaP = 0.13
  - Configured Price: $17,990

- JS21 with BladeCenter H
  - Dual-core IBM PowerPC 970MP 2.5 GHz
  - 14 2-socket blades (56 cores) in 9U chassis
    - 5676 watts / 19356 BTU for 14 blades
    - 405 watts / 1382 BTU per blade
  - Chassis: four 2900W power supplies (N+N)
  - Built-in Virtualization and vector acceleration

  - SWaP = 0.26 or 2X better than the X8400
  - Configured Price: $9,227

*For supporting information, please see next page.

Figure 8-66. Sun Blade X8400 versus JS21 Blade with BladeCenter H AM715.0

Notes:

An Open Door

- Sun is opening the door for us
  - Going from a single architecture to three new architectures
    - Niagara – new HW
    - Opteron – new HW and Solaris – 74% Linux
    - APL – new HW
  - Forced to bid Linux on Opteron to be competitive in P/P
- HP continues to be vulnerable
  - Alpha and PA-RISC clients getting to where they have written off old equipment
  - Itanium is still struggling to get market momentum – now 10% of HP's total server revenue; HP-UX down 10% YTY

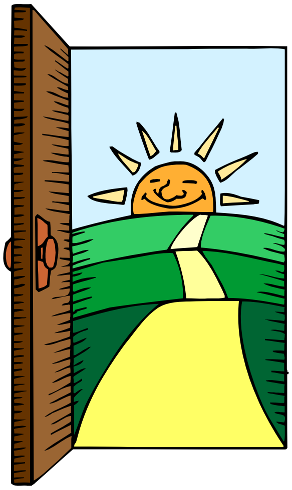

Figure 8-67. An Open Door AM715.0

Notes:

Continuing to Drive the IBM Systems Agenda

- Position high-end IBM System p as the server consolidation platform for UNIX and Linux operating systems
- Position System p servers as leadership platform for target workload / solution areas
- Refresh of System p product line with IBM POWER5+ processor and Quad-Core Module systems

Figure 8-68. Continuing to Drive the IBM Systems Agenda AM715.0

Notes:

Two-pronged July Launch Strategy

1. Achieve a more efficient data center for low TCO and operational efficiencies with IBM System p5 590/595 with POWER5+ and Advance POWER Virtualization
   - Get there with the IBM Server Consolidation Factory for System p

2. Get more for your processing dollar when moving workloads to IBM System p
   - Speed time to business value with tested solutions.
   - Introducing more new solutions
   - Featuring the IBM System p5 Express family (now at 2.1 GHz DCM and 1.65 GHz QCM)

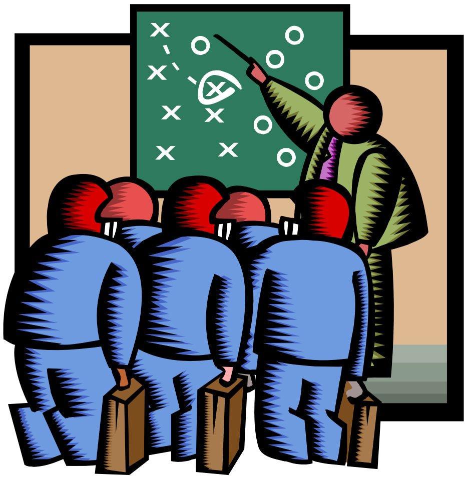

Figure 8-69. Two-pronged July Launch Strategy AM715.0

Notes:

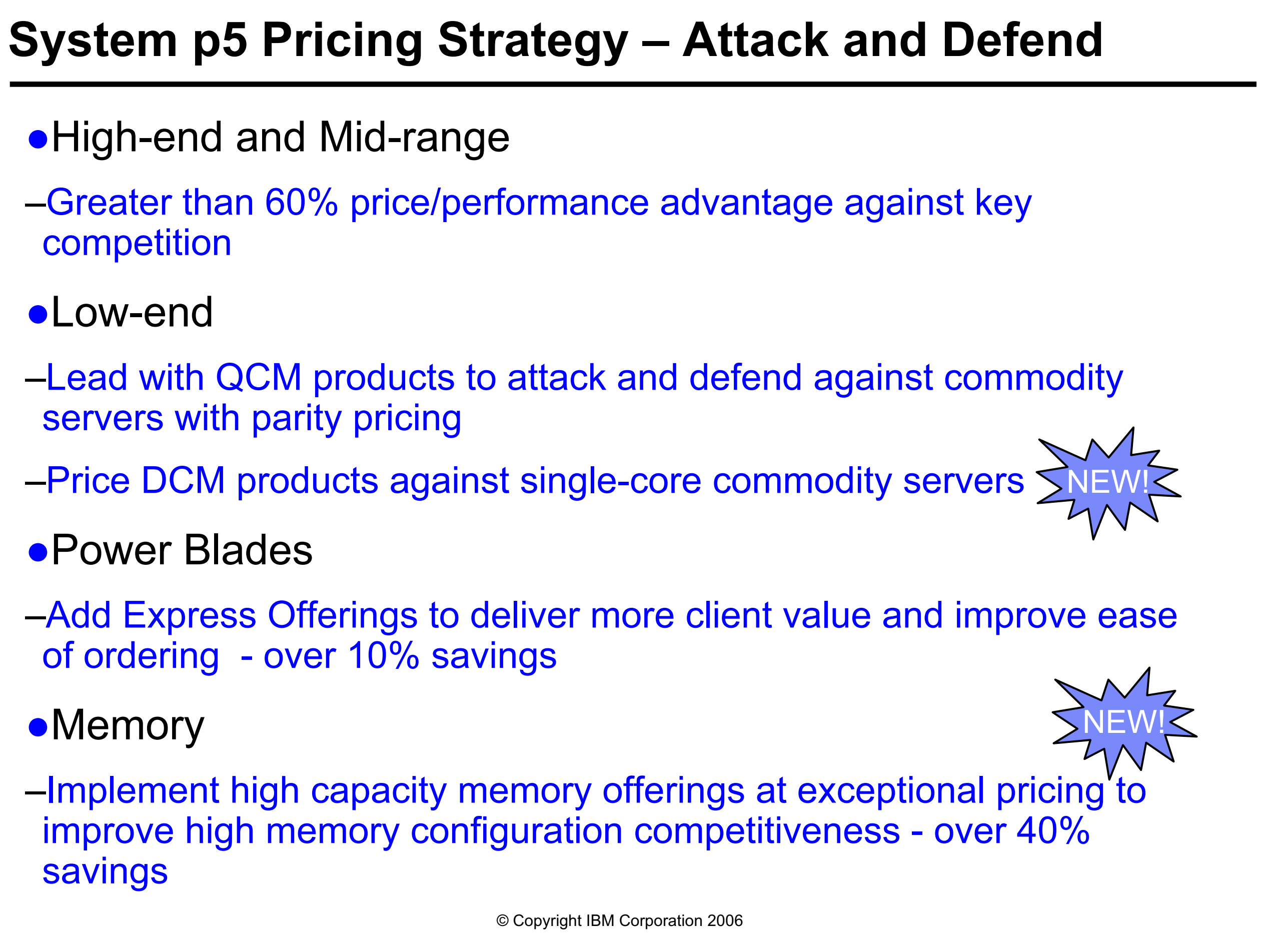

System p5 Pricing Strategy – Attack and Defend

- High-end and Mid-range
  - Greater than 60% price/performance advantage against key competition
- Low-end
  - Lead with QCM products to attack and defend against commodity servers with parity pricing
  - Price DCM products against single-core commodity servers NEW!
- Power Blades
  - Add Express Offerings to deliver more client value and improve ease of ordering - over 10% savings
- Memory NEW!
  - Implement high capacity memory offerings at exceptional pricing to improve high memory configuration competitiveness - over 40% savings

Figure 8-70. System p5 Pricing Strategy – Attack and Defend AM715.0

Notes:

IBM BladeCenter JS2x Configuration

- New System x and BladeCenter configuration capability aligning with System p
  - Allows you to spend more time on selling (US and Canada)
- Two configurator / ordering options:

1. Web-based option
   - Convenient and simple tool for less complex orders (for example, aftermarket upgrade features)
2. eConfigurator option
   - Current tool System p sellers use today – updated your eConfig profile to include BladeCenter
   - Ideally suited for build-to-order system sales
   - AIX 5L and JS2x Blade both in eConfig makes it easier to configure orders
   - Easier to save and update configurations

Figure 8-71. IBM BladeCenter JS2x Configuration AM715.0

Notes:

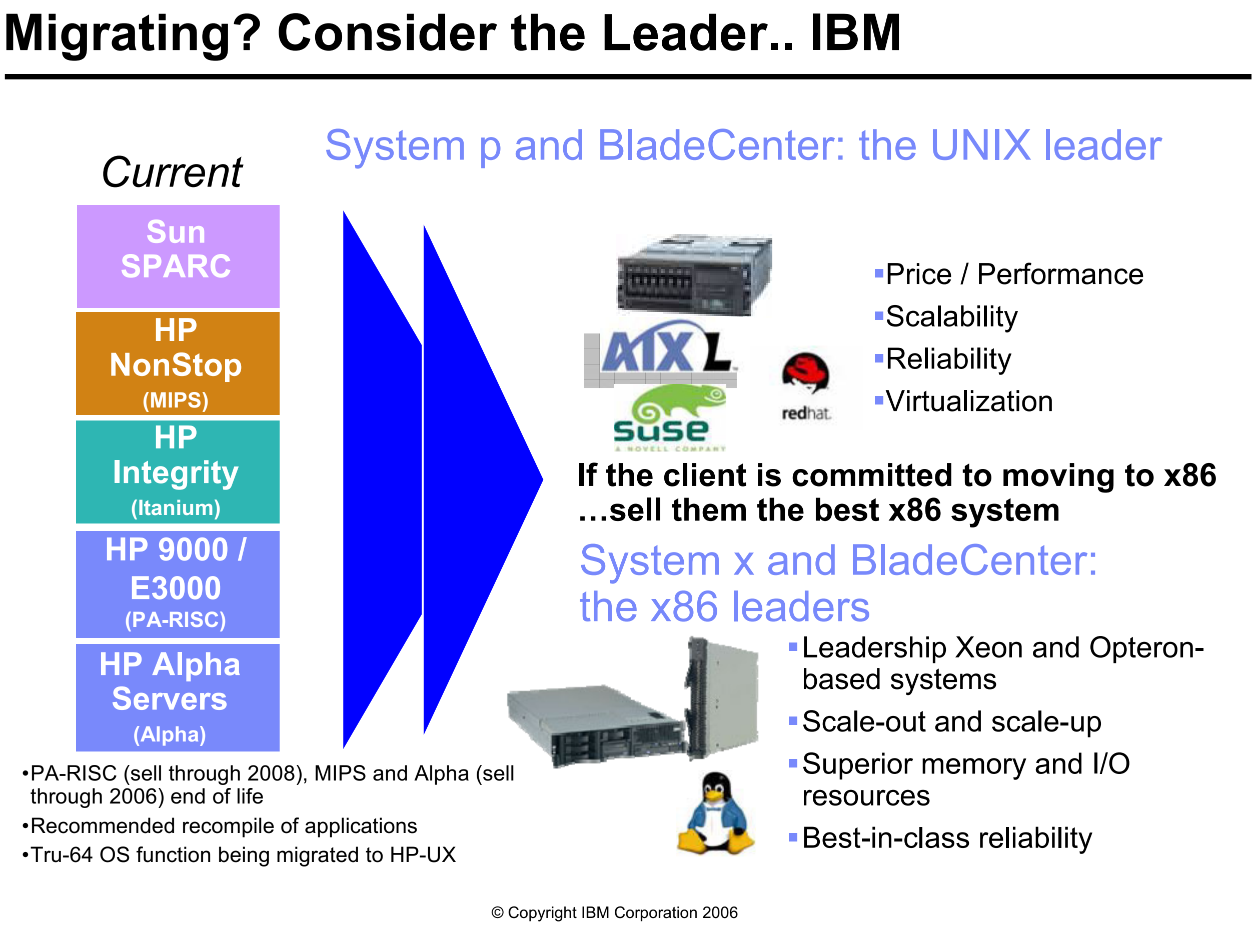

Migrating? Consider the Leader.. IBM

Current

Sun SPARC

HP NonStop (MIPS)

HP Integrity (Itanium)

HP 9000 / E3000 (PA-RISC)

HP Alpha Servers (Alpha)

- PA-RISC (sell through 2008), MIPS and Alpha (sell through 2006) end of life
- Recommended recompile of applications
- Tru-64 OS function being migrated to HP-UX

System p and BladeCenter: the UNIX leader

- Price / Performance
- Scalability
- Reliability
- Virtualization

If the client is committed to moving to x86 …sell them the best x86 system

System x and BladeCenter: the x86 leaders

- Leadership Xeon and Opteron-based systems
- Scale-out and scale-up
- Superior memory and I/O resources
- Best-in-class reliability

Figure 8-72. Migrating? Consider the Leader.. IBM AM715.0

Notes:

Performance and Benchmarks

- Commercial benchmarks are useful because they allow clients to compare performance of vendors' systems on same application
- But...
  - Benchmark workloads may not be equivalent to clients'
  - Benchmark system configurations are specific
- For comparing performance between IBM System p servers use the rperf values

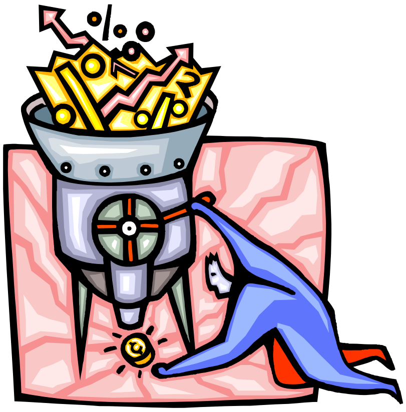

Figure 8-73. Performance and Benchmarks AM715.0

Notes:

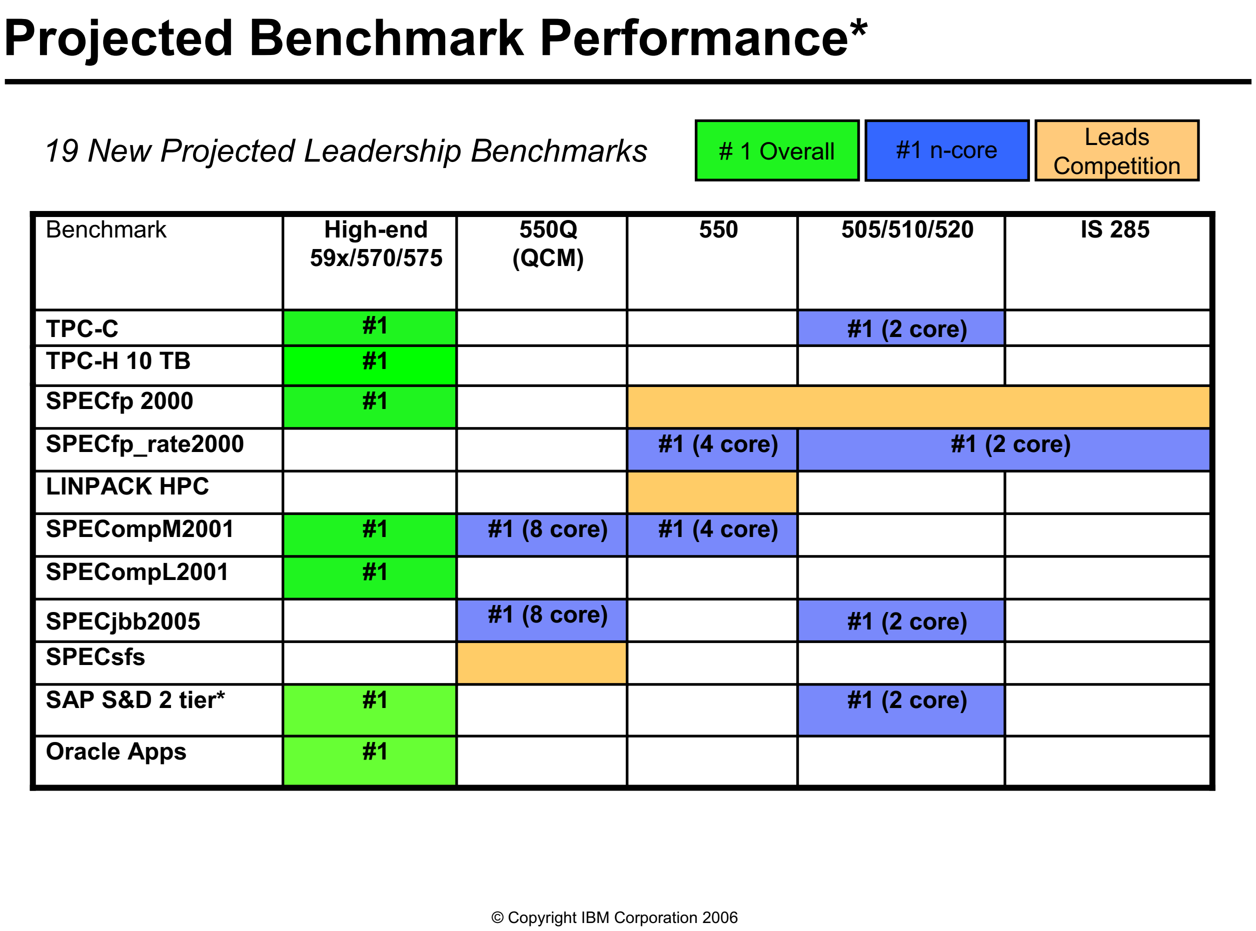

Projected Benchmark Performance*

19 New Projected Leadership Benchmarks

🟩 # 1 Overall

🟦 #1 n-core

🟧 Leads Competition

Benchmark	High-end 59x/570/575	550Q (QCM)	550	505/510/520	IS 285
TPC-C	🟩 **#1**			🟦 **#1 (2 core)**	
TPC-H 10 TB	🟩 **#1**				
SPECfp 2000	🟩 **#1**		🟧		
SPECfp_rate2000			🟦 **#1 (4 core)**	🟦 **#1 (2 core)**	
LINPACK HPC			🟧		
SPECompM2001	🟩 **#1**	🟦 **#1 (8 core)**	🟦 **#1 (4 core)**		
SPECompL2001	🟩 **#1**				
SPECjbb2005		🟦 **#1 (8 core)**		🟦 **#1 (2 core)**	
SPECsfs		🟧			
SAP S&D 2 tier*	🟩 **#1**			🟦 **#1 (2 core)**	
Oracle Apps	🟩 **#1**				

Figure 8-74. Projected Benchmark Performance* AM715.0

Notes:

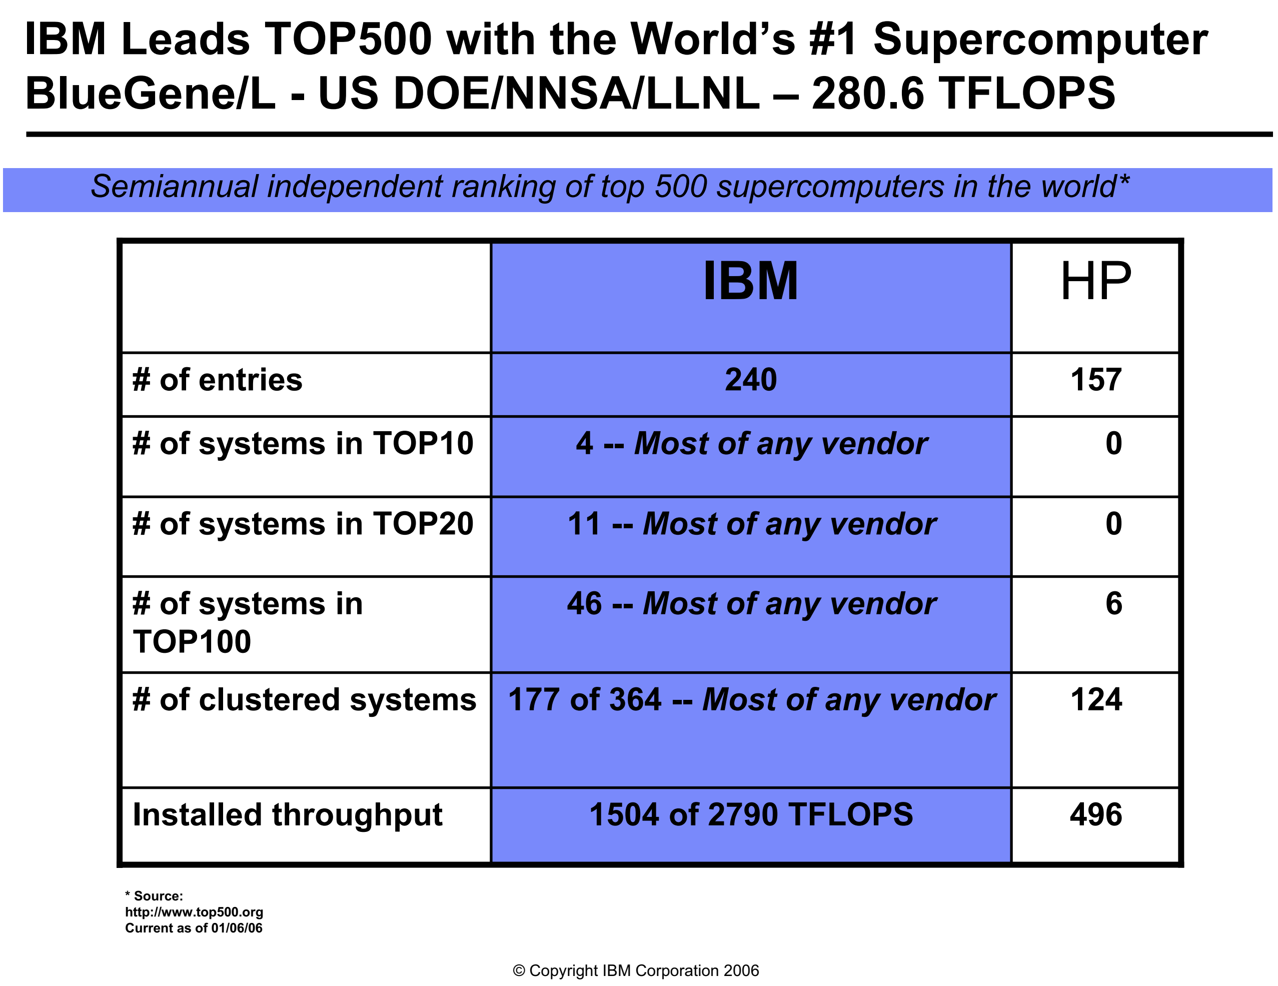

IBM Leads TOP500 with the World's #1 Supercomputer BlueGene/L - US DOE/NNSA/LLNL – 280.6 TFLOPS

*Semiannual independent ranking of top 500 supercomputers in the world**

	IBM	HP
# of entries	**240**	**157**
# of systems in TOP10	**4 -- *Most of any vendor***	**0**
# of systems in TOP20	**11 -- *Most of any vendor***	**0**
# of systems in TOP100	**46 -- *Most of any vendor***	**6**
# of clustered systems	**177 of 364 -- *Most of any vendor***	**124**
Installed throughput	**1504 of 2790 TFLOPS**	**496**

*** Source:**
http://www.top500.org
Current as of 01/06/06

Figure 8-75. IBM Leads TOP500 with the World's #1 Supercomputer BlueGene/L - US DOE/NNSA/LLNL – 280.6 TFLOPS AM715.0

Notes:

Notes on Benchmarks and Values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM System p5, eServer p5, pSeries, OpenPower, RS/6000 and BladeCenter Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, AIX Version 4.3 or AIX 5L were used. All other systems used previous versions of AIX. The SPEC CPU2000, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C Enterprise Edition V7.0 for AIX, XL C/C++ Enterprise Edition V7.0 for AIX, XL FORTRAN Enterprise Edition V9.1 for AIX, XL C/C++ Advanced Edition V7.0 for Linux, and XL FORTRAN Advanced Edition V9.1 for Linux. The SPEC CPU95 (retired in 2000) tests used preprocessors, KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX, MASS for AIX and Kazushige Goto's BLAS Library for Linux were also used in some benchmarks.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
NotesBench	http://www.notesbench.org
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle Applications	http://www.oracle.com/apps_benchmark/
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	http://www.siebel.com/crm/performance_benchmark/index.shtm
Baan	http://www.ssaglobal.com
Microsoft Exchange	http://www.microsoft.com/exchange/evaluation/performance/default.asp
Veritest	http://www.veritest.com/clients/reports
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised April 27, 2006

Figure 8-76. Notes on Benchmarks and Values AM715.0

Notes:

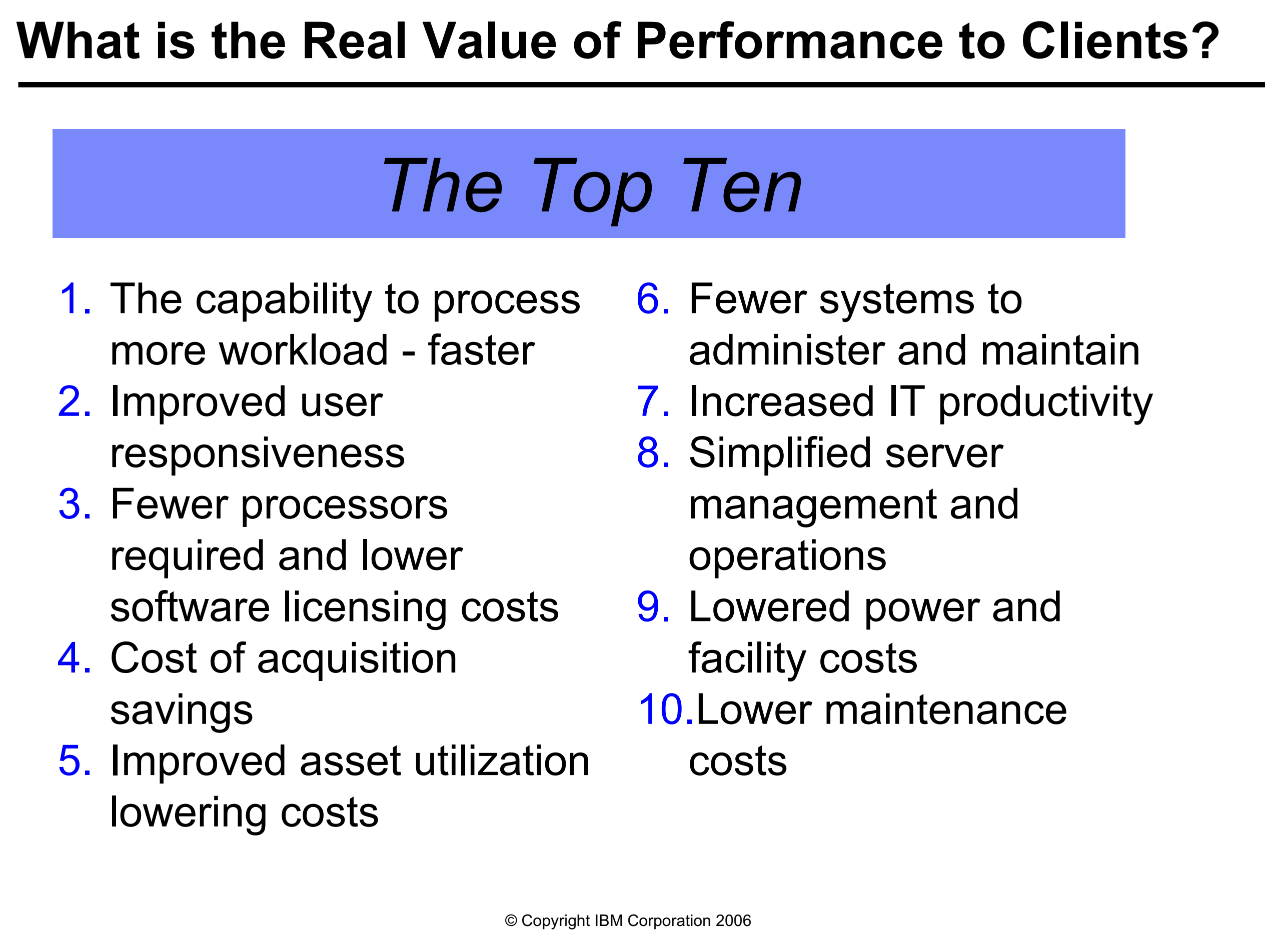

What is the Real Value of Performance to Clients?

The Top Ten

1. The capability to process more workload - faster
2. Improved user responsiveness
3. Fewer processors required and lower software licensing costs
4. Cost of acquisition savings
5. Improved asset utilization lowering costs
6. Fewer systems to administer and maintain
7. Increased IT productivity
8. Simplified server management and operations
9. Lowered power and facility costs
10. Lower maintenance costs

Figure 8-77. What is the Real Value of Performance to Clients? AM715.0

Notes:

POWER5 RAS Features

Primary POWER5 RAS Features

- ✓ First Failure Data Capture
- ✓ DDR Chipkill memory
- ✓ Bit-steering/redundant memory
- ✓ Memory soft scrubbing
- ✓ Redundant power, fans
- ✓ Dynamic Processor Deallocation
- ✓ Dynamic processor sparing
- ✓ ECC memory
- ✓ Persistent memory deallocation
- ✓ Hot-plug PCI slots, fans, power
- ✓ Internal light path diagnostics
- ✓ Hot-swappable disk bays
- ✓ I/O error handling extended beyond base PCI adapter
- ✓ ECC extended to inter-chip connections for the fabric/processor buses

2005 RAS Enhancements

- **Service Processor Failover:**
  - p5-570, p5-590 and p5-595
  - New service processor option for p5-570
    - ✓ Required for “Failover” support
- **Dynamic Firmware Maintenance:**
  - Apply firmware fixes without system disruption
  - Fixes only, not new functionality
- **Hot I/O Drawer Add**
  - Install remote I/O drawers without system disruption

HMC required to enable these functions..

Figure 8-78. POWER5 RAS Features AM715.0

Notes:

Virtualization in the Future

- Effective hypervisors will become pervasive
  - EMC VMware, Microsoft Virtual Server, Xen open source hypervisor, HP VM hypervisor…
- Hypervisor support will be provided by Intel and AMD
  - Intel (VT-x, Oct. 2005; VT-i, 2006) and AMD (Pacifica, 1H06)
- Self-virtualizing I/O adapters will become available, enabling direct I/O adapter sharing
  - Standard I/O adapters that are self-virtualizing are likely to become available in 2-3 years

Figure 8-79. Virtualization in the Future AM715.0

Notes:

Virtualization in the Future

- Operating systems will become hypervisor aware / exploitive.
  - Linux today supports paravirtualization for POWER, and will soon do so for IA as well.
  - Microsoft has stated that future Windows versions will have MS VS “enlightenment.”
- Hypervisors will become integrated as firmware on high volume servers in 1-2 years.
  - This will avoid the need for installing an add-on and provide platform differentiation.

Figure 8-80. Virtualization in the Future AM715.0

Notes:

Virtualization in the Future

- Hypervisors will become bundled as software with operating systems.
  - Linux distros will include Xen in 4Q05-1Q06
  - Microsoft will introduce a Type 1 form of MS VS and integrate it with Windows (2007?)
- Application containers will gain use as an efficient method for consolidating many applications per OS.
  - Sun Solaris containers, SWsoft Virtuozzo containers for Windows/Linux/UNIX, HP-UX Secure Resource Partitions, AIX Corrals leveraging Meiosys, VERITAS containers with Live Migration, Softricity SystemGuard/SoftGrid

Figure 8-81. Virtualization in the Future AM715.0

Notes:

Available Partitioning Approaches

- Hardware Partitioning
  - Strengths:
    - Hardware fault isolation
    - Least overhead
  - Weaknesses:
    - Poor granularity
    - No resource sharing
    - Limited flexibility
- Firmware or Hypervisor
  - Strengths:
    - Best granularity
    - Most resource sharing
    - Greatest flexibility
  - Weaknesses:
    - Overhead
    - Fault isolation

Figure 8-82. Available Partitioning Approaches AM715.0

Notes:

Partitioning Leadership

Capability	POWER5	Other UNIX Hypervisors	H/W Partitioning
Partition Granularity	1/10th	1	4 / 8
Memory Granularity	128 MB	2 GB	2 GB
Max # of Partitions	254	128	128 / 144
I/O Adapter Sharing	Yes	No	No
Multi-OS Support	3	1	4 / 1

Figure 8-83. Partitioning Leadership

AM715.0

Notes:

Business Value of Advanced POWER Virtualization

- Reduce costs (increased asset utilization)
- Lower management costs (fewer systems to manage)
- Increase IT productivity (through simplified server management and operations)

Figure 8-84. Business Value of Advanced POWER Virtualization AM715.0

Notes:

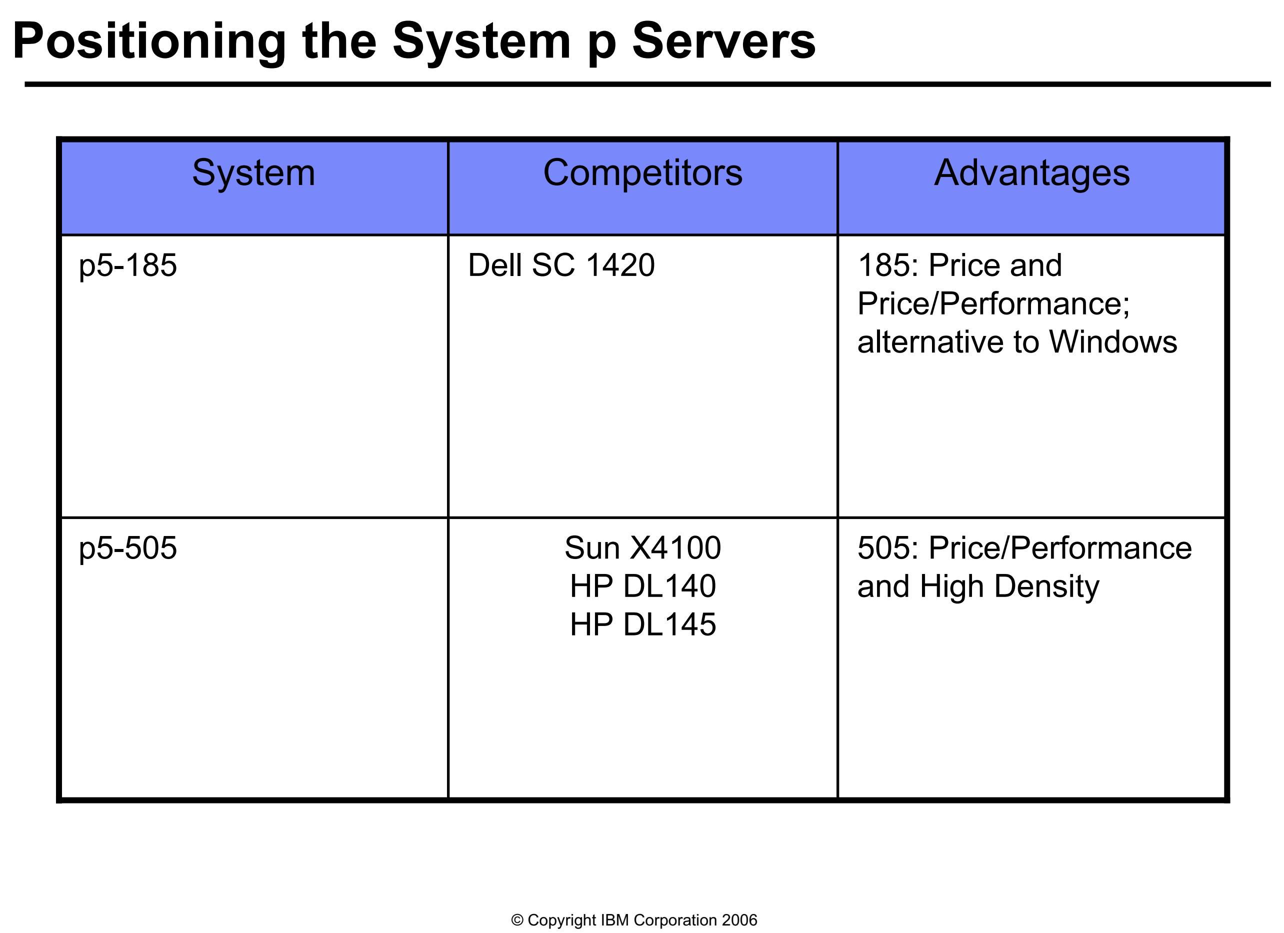

Positioning the System p Servers

System	Competitors	Advantages
p5-185	Dell SC 1420	185: Price and Price/Performance; alternative to Windows
p5-505	Sun X4100 HP DL140 HP DL145	505: Price/Performance and High Density

Figure 8-85. Positioning the System p Servers

AM715.0

Notes:

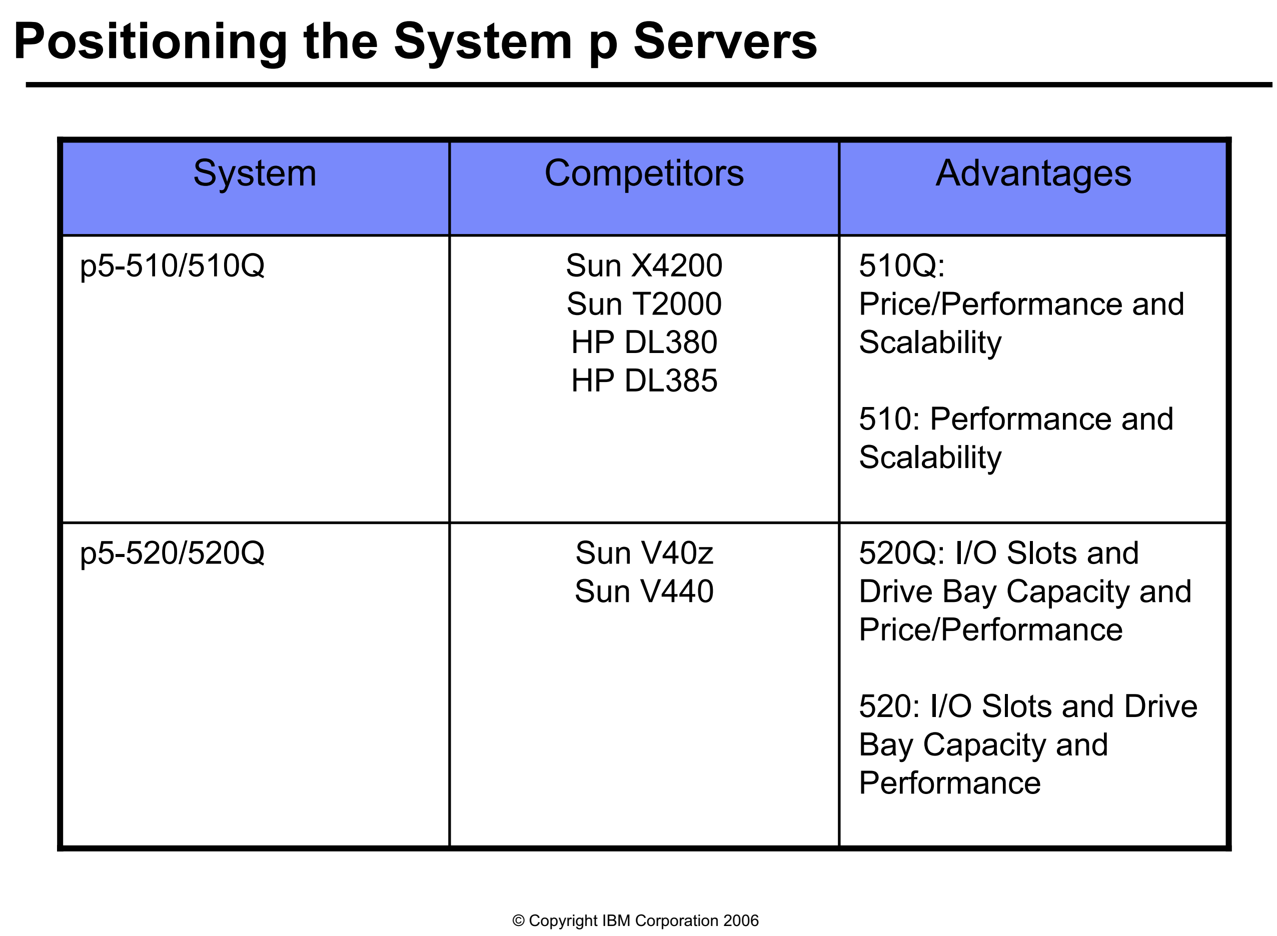

Positioning the System p Servers

System	Competitors	Advantages
p5-510/510Q	Sun X4200 Sun T2000 HP DL380 HP DL385	510Q: Price/Performance and Scalability 510: Performance and Scalability
p5-520/520Q	Sun V40z Sun V440	520Q: I/O Slots and Drive Bay Capacity and Price/Performance 520: I/O Slots and Drive Bay Capacity and Performance

Figure 8-86. Positioning the System p Servers AM715.0

Notes:

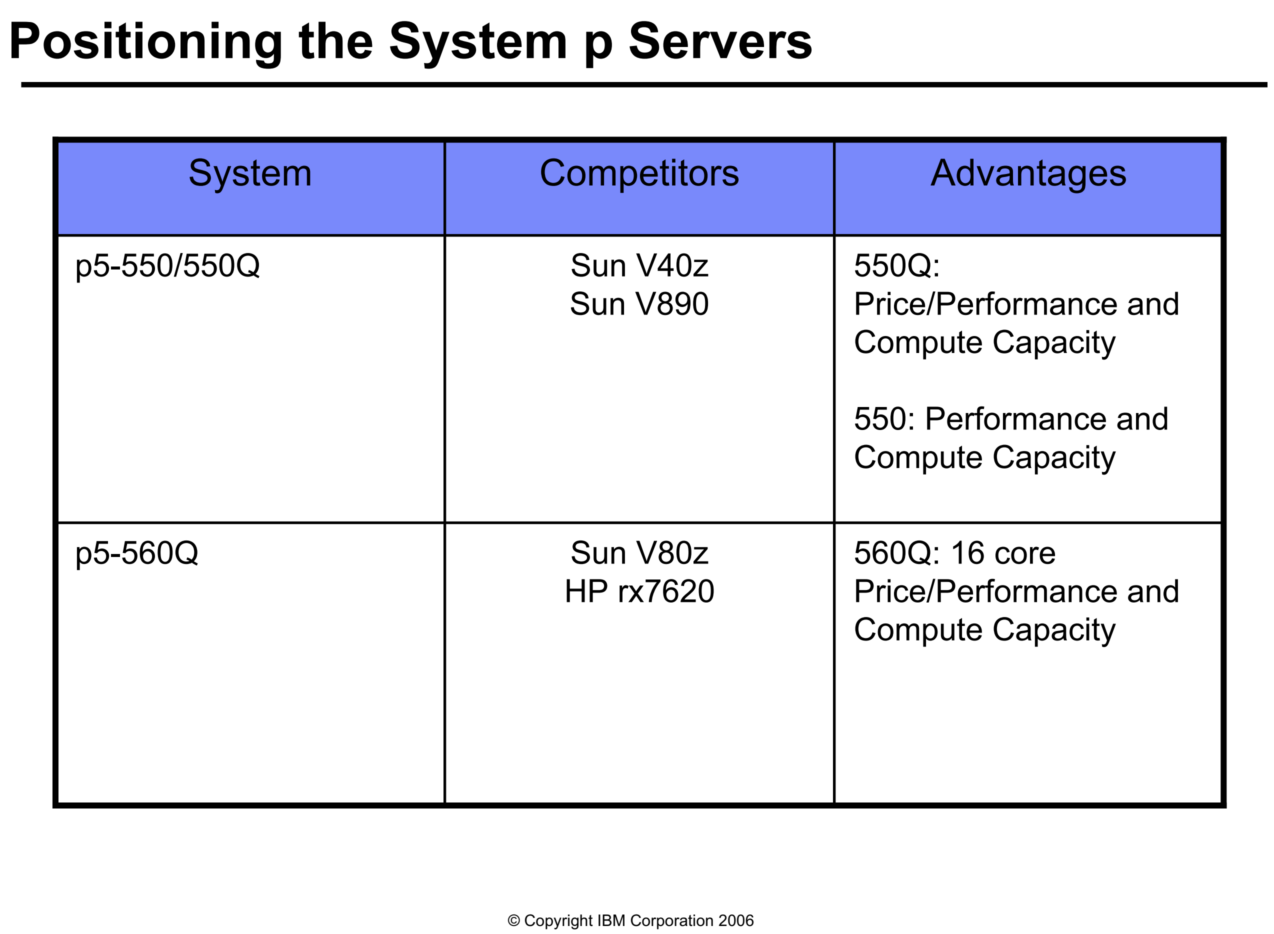

Positioning the System p Servers

System	Competitors	Advantages
p5-550/550Q	Sun V40z Sun V890	550Q: Price/Performance and Compute Capacity 550: Performance and Compute Capacity
p5-560Q	Sun V80z HP rx7620	560Q: 16 core Price/Performance and Compute Capacity

Figure 8-87. Positioning the System p Servers AM715.0

Notes:

p5-570 Positioning

Competitors	Target Sectors	Differentiation
	1.65 GHz	
•HP rx7620	**•Financial Services**	**•Virtualization**
•HP rx8620	**•Industrial**	**•Performance**
•HP rp7420	**•Distributors**	**•Price/Performance**
•HP rp8420	**•Public Sector**	**•Cost of Ownership**
•Sun E2900	**•Retail**	
•Sun E6900	**•Communications**	
•Sun V1280	**•HPC**	
	1.9 GHz	
•Premium Offering	**•Financial Services**	**•Performance**
	•Industrial	**•Virtualization**
	•Distributors	**•Price/Performance**
	•Public Sector	**•Cost of Ownership**
	•Retail	
	•Communications	
	•HPC	

- 1.65 GHz targeted at the mid-range UNIX market
- Competes with bulk of UNIX mid-range systems
- Wide range of differentiation

- 1.9 GHz targeted at Enterprise UNIX market
- Premium mid-range offering with leadership performance*
- Performance the primary differentiator

* Performance claim based on POWER5 versus Best of Breed Benchmark Chart

Figure 8-88. p5-570 Positioning AM711.0

Notes::

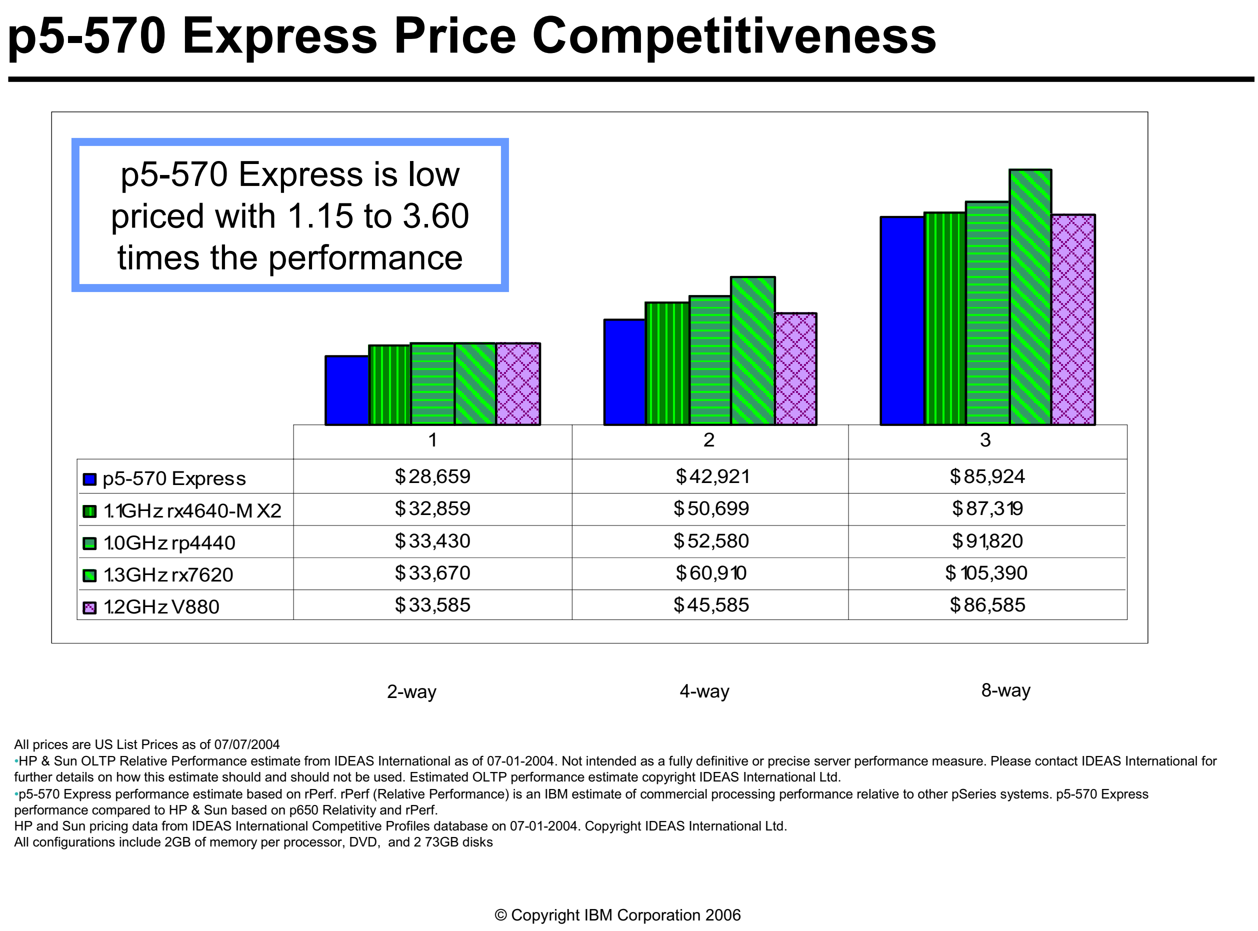

p5-570 Express Price Competitiveness

p5-570 Express is low priced with 1.15 to 3.60 times the performance

	1	2	3
■ p5-570 Express	$28,659	$42,921	$85,924
▥ 1.1GHz rx4640-M X2	$32,859	$50,699	$87,319
▤ 1.0GHz rp4440	$33,430	$52,580	$91,820
▨ 1.3GHz rx7620	$33,670	$60,910	$105,390
▩ 1.2GHz V880	$33,585	$45,585	$86,585

All prices are US List Prices as of 07/07/2004

- HP & Sun OLTP Relative Performance estimate from IDEAS International as of 07-01-2004. Not intended as a fully definitive or precise server performance measure. Please contact IDEAS International for further details on how this estimate should and should not be used. Estimated OLTP performance estimate copyright IDEAS International Ltd.
- p5-570 Express performance estimate based on rPerf. rPerf (Relative Performance) is an IBM estimate of commercial processing performance relative to other pSeries systems. p5-570 Express performance compared to HP & Sun based on p650 Relativity and rPerf.

HP and Sun pricing data from IDEAS International Competitive Profiles database on 07-01-2004. Copyright IDEAS International Ltd.
All configurations include 2GB of memory per processor, DVD, and 2 73GB disks

Figure 8-89. p5-570 Express Price Competitiveness AM715.0

Notes:

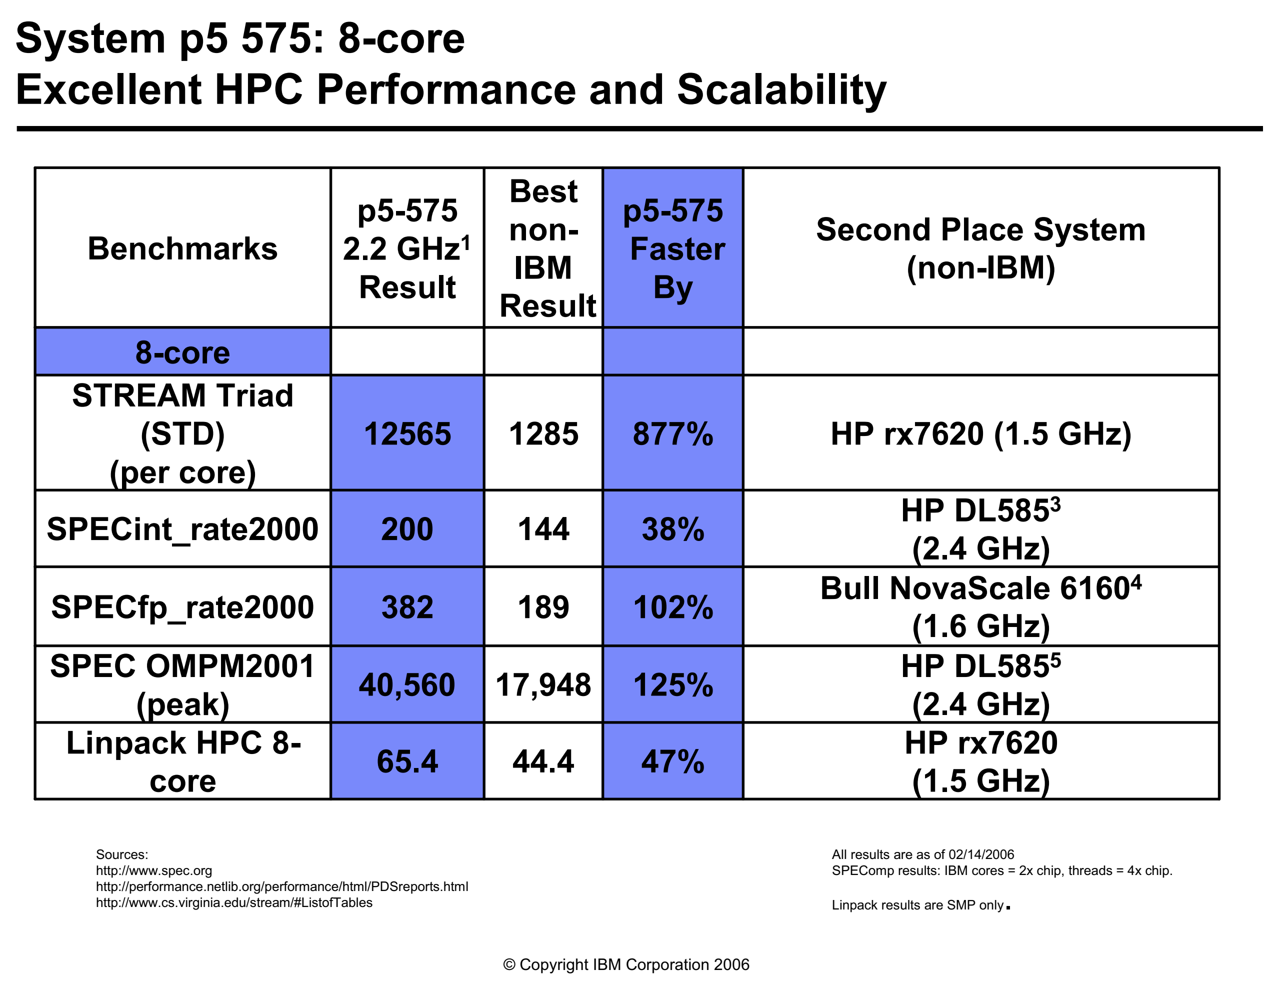

System p5 575: 8-core
Excellent HPC Performance and Scalability

Benchmarks	p5-575 2.2 GHz[1] Result	Best non-IBM Result	p5-575 Faster By	Second Place System (non-IBM)
8-core				
STREAM Triad (STD) (per core)	12565	1285	877%	HP rx7620 (1.5 GHz)
SPECint_rate2000	200	144	38%	HP DL585[3] (2.4 GHz)
SPECfp_rate2000	382	189	102%	Bull NovaScale 6160[4] (1.6 GHz)
SPEC OMPM2001 (peak)	40,560	17,948	125%	HP DL585[5] (2.4 GHz)
Linpack HPC 8-core	65.4	44.4	47%	HP rx7620 (1.5 GHz)

Sources:
http://www.spec.org
http://performance.netlib.org/performance/html/PDSreports.html
http://www.cs.virginia.edu/stream/#ListofTables

All results are as of 02/14/2006
SPEComp results: IBM cores = 2x chip, threads = 4x chip.

Linpack results are SMP only.

Figure 8-90. System p5 575:8-core Excellent HPC Performance and Scalability AM715.0

Notes:

System p5 575: 16-core
Excellent HPC Performance and Scalability

16-core	p5-575 1.9 GHz[2] Result	Best non-IBM result	p5-575 Faster By	Second Place System (non-IBM)
STREAM Triad (STD) (per core)	5378	1316	308%	HP Integrity SD (1.6 GHz)
SPECint_rate2000	314	266	18%	HP rx8620-32 (1.6 GHz)[6]
SPECfp_rate2000	571	373	53%	Bull NovaScale 6320[7] (1.6 GHz)
SPEC OMPM2001 (peak)	56,211	25,932	116%	HP AlphaServer GS1280[8] (1.3 GHz)
Linpack HPC 16-core	111.4	88.8	25%	HP rx8620 (1.5 GHz)

Figure 8-91. System p5 575:16-core Excellent HPC Performance and Scalability AM715.0

Notes:

IBM Data Warehousing Balanced Configuration Unit

Standardized, assured, data warehouse solution building block

- Preconfigured building blocks:
  - 8-core 1.9 GHz p5-575
  - DS4500, software and support*
- Fully integrated and tested package
- Uses standard IBM hardware and software components
- Can be used with BI software applications already qualified on those platforms such as Siebel, SAS, SAP and others**

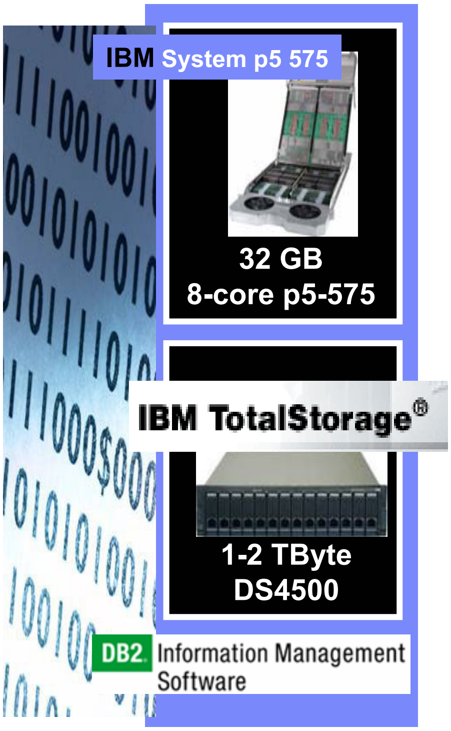

*(1) 32GB p5-575 per (1-2) TB of data
**http://www.ibm.com/software/data/db2bi/bcu/
Current as of 09/27/2005

Figure 8-92. IBM Data Warehousing Balanced Configuration Unit AM715.0

Notes:

p5-595: Consolidation on a Whole New Level

The capacity to handle the workload of
32 systems and 256 processors -- with room to spare

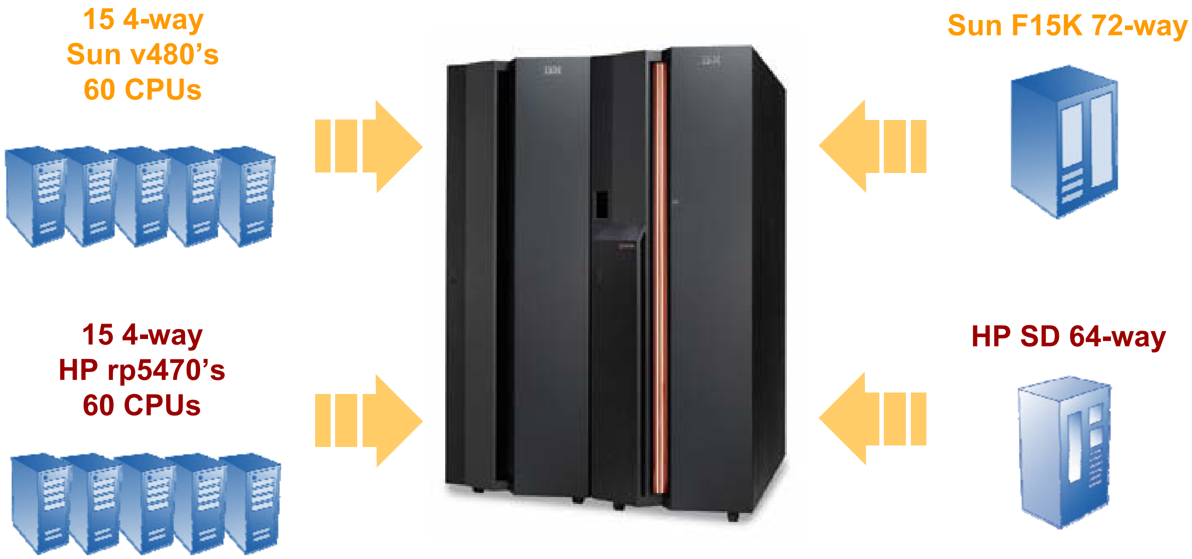

15 4-way
Sun v480's
60 CPUs

15 4-way
HP rp5470's
60 CPUs

Sun F15K 72-way

HP SD 64-way

Consolidation example based on performance assessments available through IDEAS International Competitive Profiles database as of 10/13/04. Copyright IDEAS International LTD. www.ideasinternational.com

Figure 8-93. p5-595: Consolidation on a Whole New Level* AM715.0

Notes:

Checkpoint Questions (1 of 2)

1. True or False. Commercial benchmarks results can be easily compared to the client’s workload.

2. IBM System p has a competitive advantage in which of the following?
    a. Performance
    b. RAS
    c. Virtualization
    d. All of the above

3. Which of the following is an advantage of logical partitioning over physical partitioning?
    a. Granularity
    b. Resource sharing
    c. Flexibility
    d. All of the above

Figure 8-94. Checkpoint Questions (1 of 2) AM715.0

Notes:

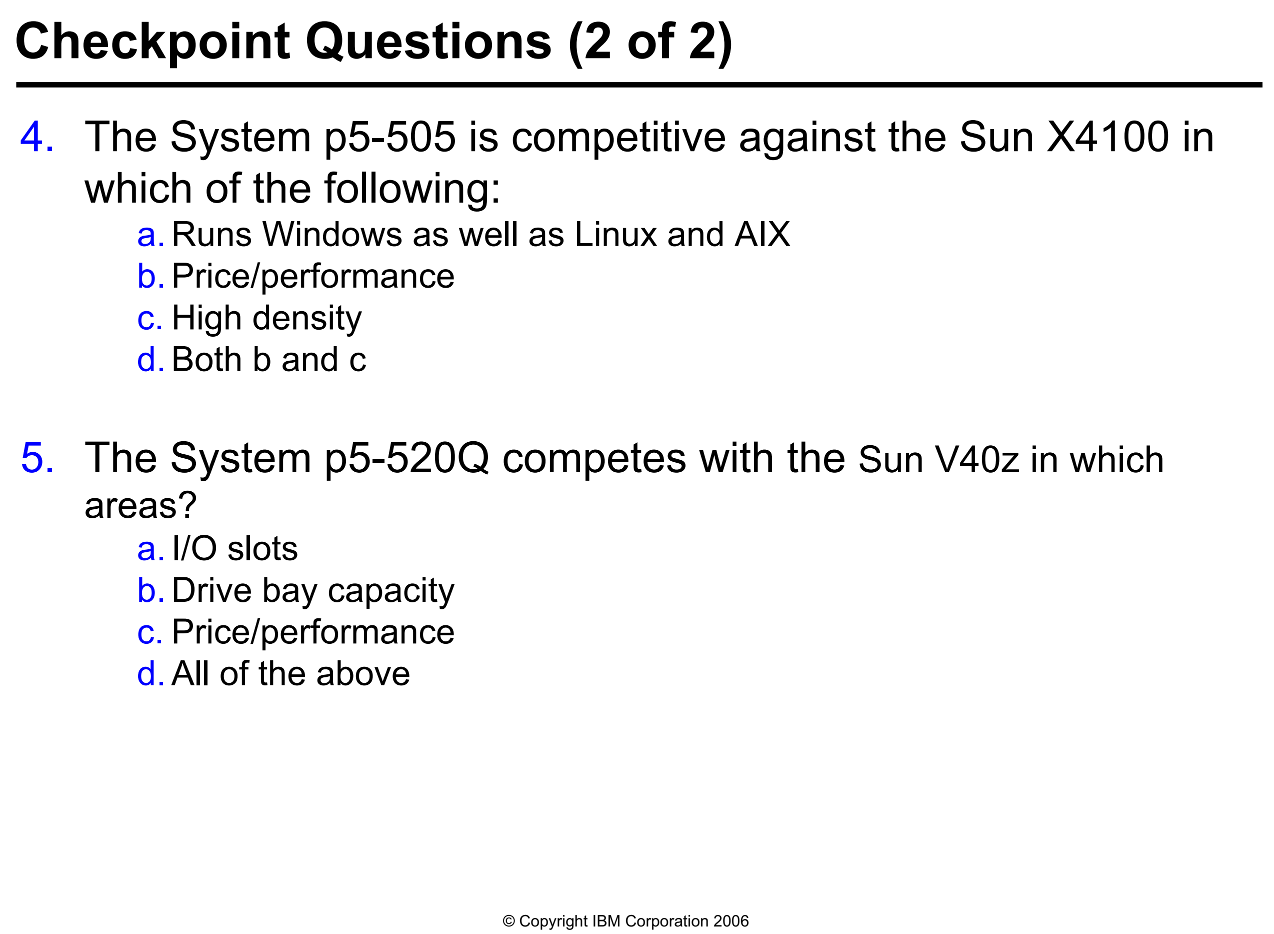

Checkpoint Questions (2 of 2)

4. The System p5-505 is competitive against the Sun X4100 in which of the following:
    - a. Runs Windows as well as Linux and AIX
    - b. Price/performance
    - c. High density
    - d. Both b and c

5. The System p5-520Q competes with the Sun V40z in which areas?
    - a. I/O slots
    - b. Drive bay capacity
    - c. Price/performance
    - d. All of the above

Figure 8-95. Checkpoint Questions (2 of 2) AM715.0

Notes:

Unit Summary

Having completed this unit, you should be able to:

- Explain how our features and capabilities provide competitive advantages.
- Position the System p products against the competition.

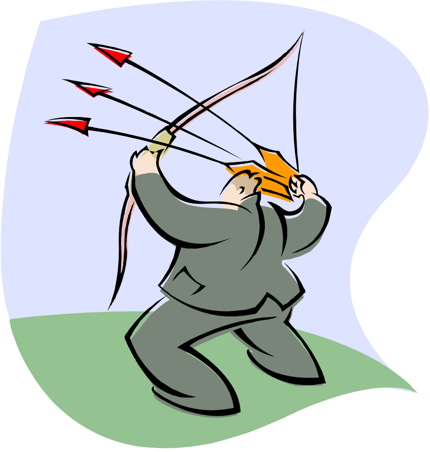

Figure 8-96. Unit Summary AM715.0

Notes:

Unit 9. Sizing the Solution

What This Unit Is About

This unit describes the different methods used to size systems.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Understand the constraints placed on accurate sizing
- Describe ‘rules of thumb’ for realistic sizing
- Explain the use of the Balanced System Guideline tool
- Explain the use of the System Planning tool
- Locate other tools for sizing systems for particular applications

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

On completion of this unit, you should be able to:

- Understand the constraints placed on accurate sizing.
- Describe 'rules of thumb' for realistic sizing.
- Explain the use of the Balanced System Guideline tool.
- Explain the use of the System Planning tool.
- Locate other tools for sizing systems for particular applications.

Figure 9-1. Unit Objectives AM711.0

Notes:

Sizing and Solution Design

- Sizing is one aspect of solution design
- Is it possible to have an accurately sized solution and still not be successful?
- How do the other aspects of solution design affect the sizing process?
  - Business goals, targets and aspirations?
  - Application software choices and middleware selection?
  - Infrastructure and organization of the IT department?
  - Communications and interfaces with other systems?

Figure 9-2. Sizing and Solution Design AM711.0

Notes:

As technical specialists we are focused on providing the technical sizing specification for the solution. Sometimes the technical part of the solution is easiest, in that we feel more in control of the information and can manipulate data in abstract ways that yield results that seem more tangible. But the solution, in order to be considered as successful by the client, must meet other criteria and respond to other pressures than whether the workload can run in an acceptable way.

- System design and requirements
- Defining business goals, targets, and aspirations
- Building a business model and data model
- Re-engineering current processes
- Application software choices and middleware selection
- Infrastructure and organization of the information technology (IT) department
- Communications and interfaces with other systems

From these activities, the data required for sizing should evolve.

Sizing Concepts

“Properly sizing pSeries servers can be difficult since every client environment is unique. Usually there is not enough information to make the best decision. Therefore you must take a realistic approach. If you make proper assumptions you can reach an adequate solution.”

IBM eServer pSeries Sizing and Capacity Planning Redbook.

Figure 9-3. Sizing Concepts AM711.0

Notes:

Part of the problem for the technical specialist is the lack of information on which to base a technical solution design. The sizing experts realize this and work on the basis of making a ‘best decision' based on information and proper assumptions.

The Sizing Model

- Data entered into model or spreadsheet
- Solution requirements are determined
  - rPerf value for processor power rating
  - Memory sizing
  - Number and size of disk drives
  - Adapters for networking and storage
- Creating the model is the key to good sizing
- Model works from sizing of data to hardware configuration

Figure 9-4. The Sizing Model AM711.0

Notes:

The person who is responsible for sizing must enter, into a modelling tool or spreadsheet, the available data on workloads, the number of users, transactions, and data volumes. From this information and data about pSeries servers, the application, benchmarks, hardware, and software requirements are determined.

These are usually expressed in a processor power rating (such as rPerfs, SPECInt, and so forth), the size of memory, number and size of disk drives, along with adapters for networking and storage.

Creating this model is key to good sizing. The model works from the sizing of data to the hardware configuration.

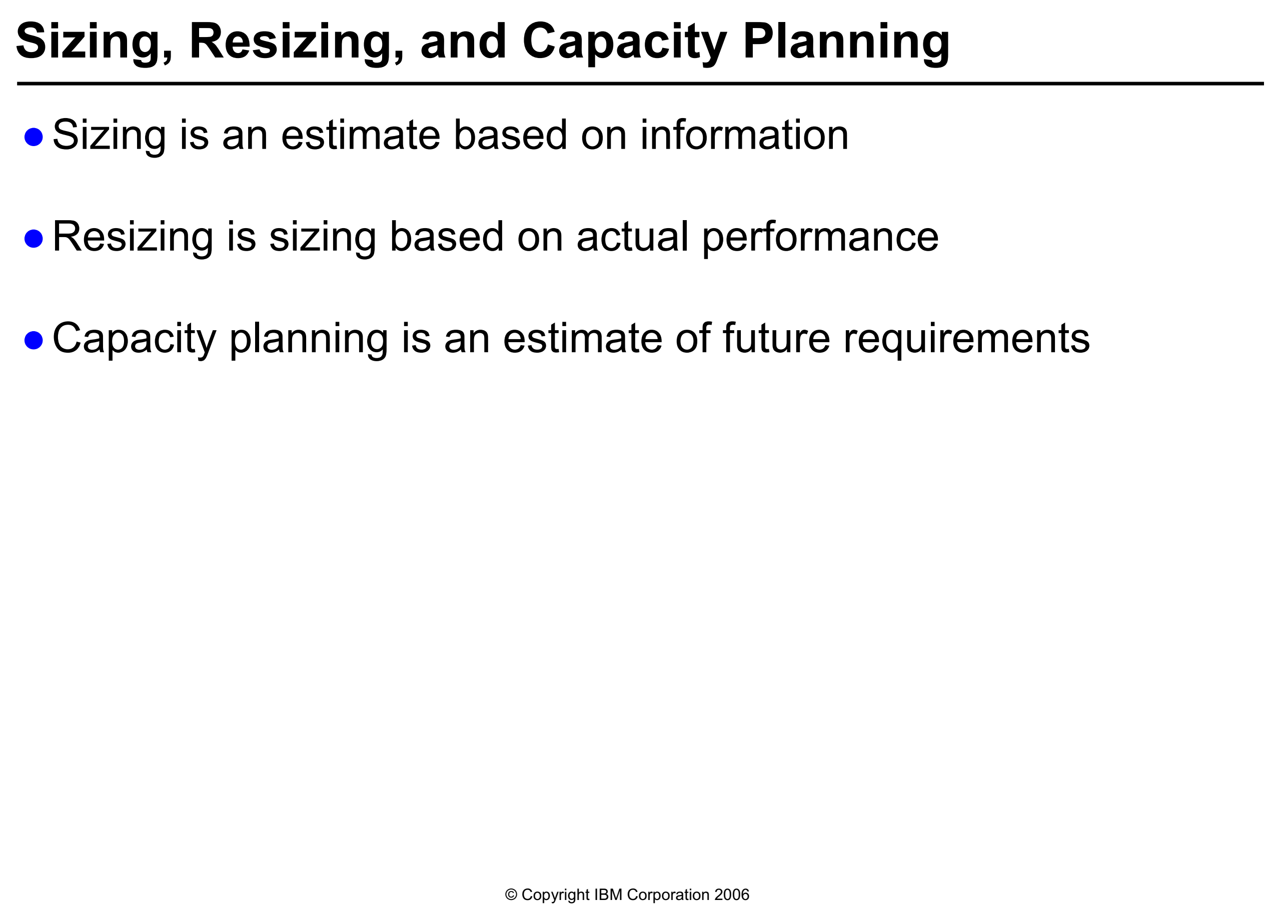

Sizing, Resizing, and Capacity Planning

- Sizing is an estimate based on information
- Resizing is sizing based on actual performance
- Capacity planning is an estimate of future requirements

Figure 9-5. Sizing, Resizing, and Capacity Planning AM711.0

Notes:

When a new server is being sized the sizing is necessarily based on estimates. The resizing process has the advantage, from the point of view of increased information, of being based on an existing system. Capacity planning is also subject to errors of estimation since it attempts to guess the future. For any of these sizing tasks one of the key determinants of accuracy is the amount and quality of information available on which to base the sizing estimate.

The Sizing Problem

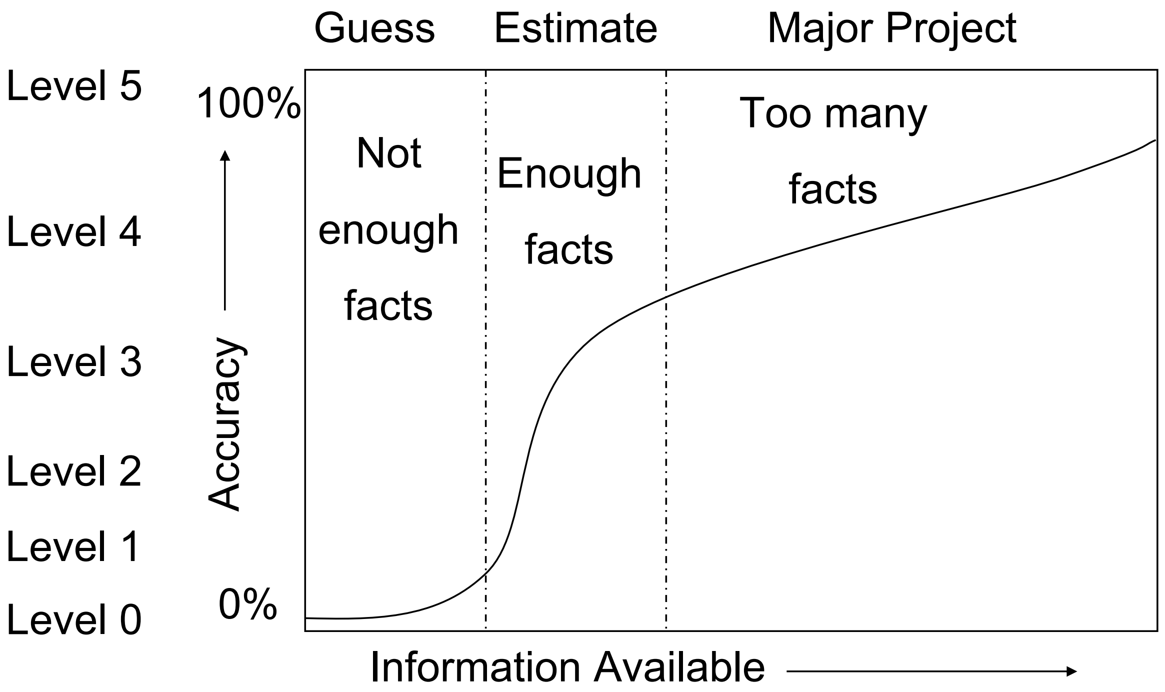

Figure 9-6. The Sizing Problem AM711.0

Notes:

Sizing is based on information. Too little and the sizing is a guess, too much and the sizing might be accurate but may take too long to fit other requirements. Ideally we are looking for enough information on which to base a realistic estimate with known assumptions.

Sizing Levels

- Level: 0
- Accuracy: 500%
- Description: Pure guess work
- Input Data: None or nearly none
- Time for Sizing: 1 second

- Level: 1
- Accuracy: 100%
- Description: Guess
- Input Data: Segmentation only
- Time for Sizing: 10 minutes

Figure 9-7. Sizing Levels AM711.0

Notes:

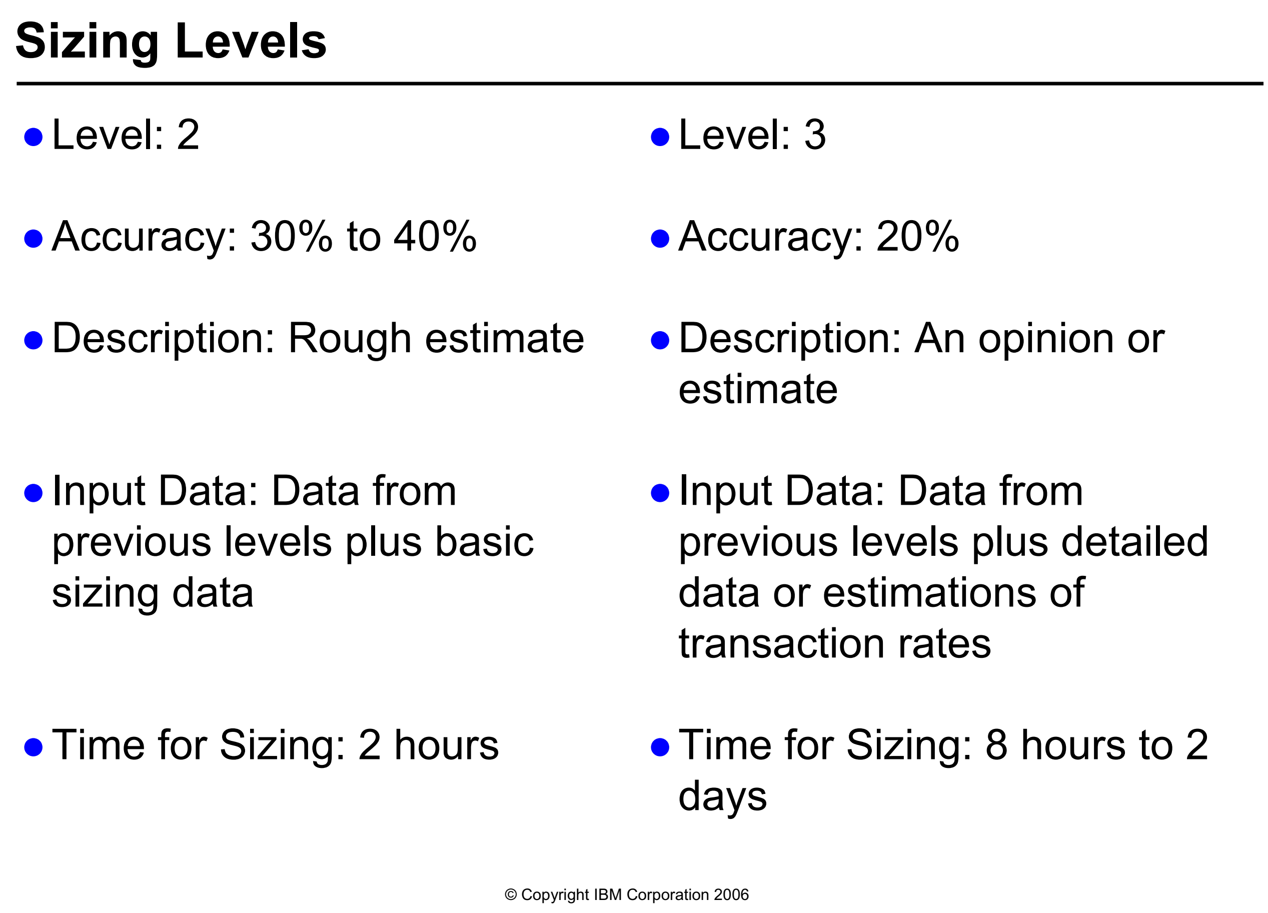

Sizing Levels

- Level: 2
- Accuracy: 30% to 40%
- Description: Rough estimate
- Input Data: Data from previous levels plus basic sizing data
- Time for Sizing: 2 hours

- Level: 3
- Accuracy: 20%
- Description: An opinion or estimate
- Input Data: Data from previous levels plus detailed data or estimations of transaction rates
- Time for Sizing: 8 hours to 2 days

Figure 9-8. Sizing Levels AM711.0

Notes:

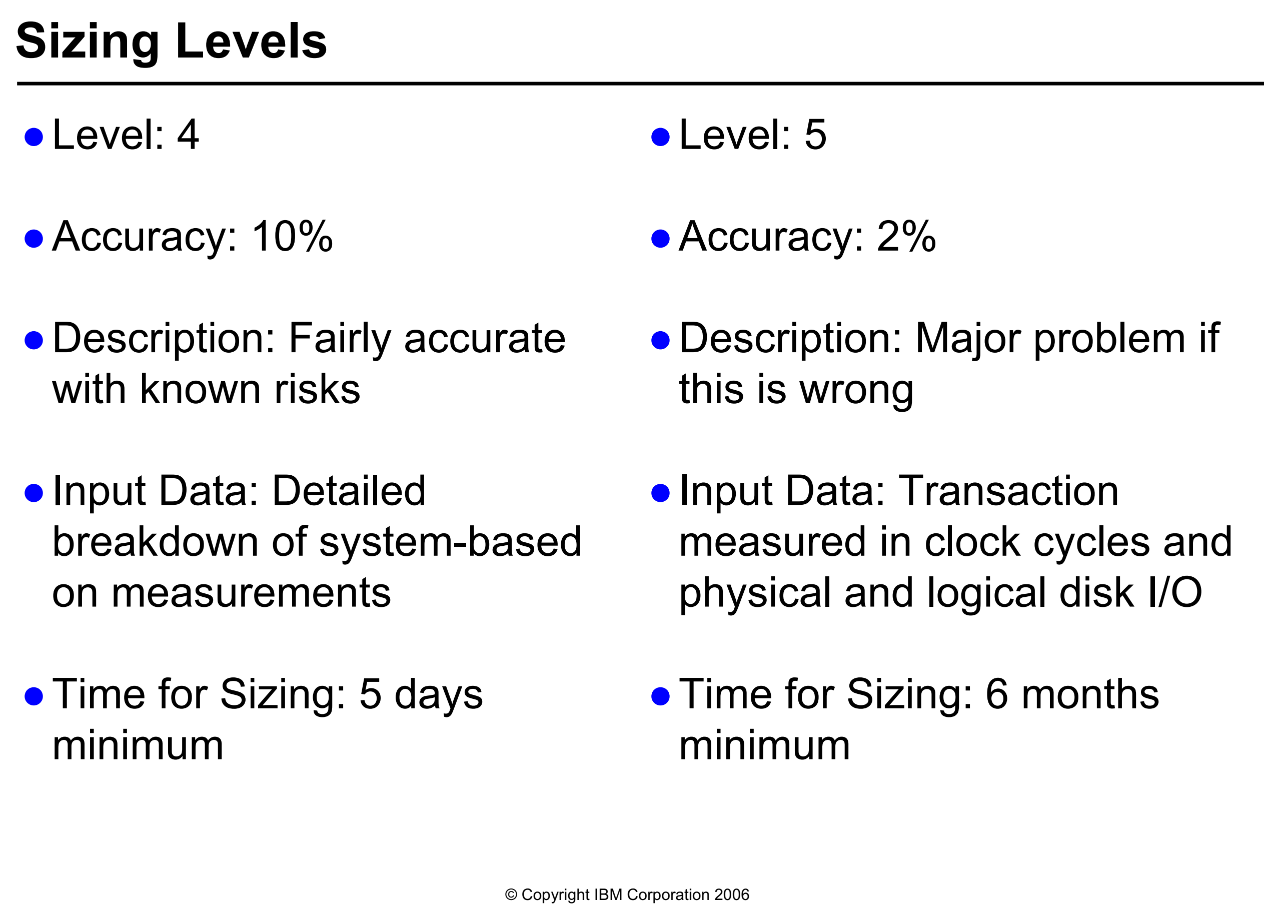

Sizing Levels

- Level: 4
- Accuracy: 10%
- Description: Fairly accurate with known risks
- Input Data: Detailed breakdown of system-based on measurements
- Time for Sizing: 5 days minimum

- Level: 5
- Accuracy: 2%
- Description: Major problem if this is wrong
- Input Data: Transaction measured in clock cycles and physical and logical disk I/O
- Time for Sizing: 6 months minimum

Figure 9-9. Sizing Levels

AM711.0

Notes:

Sizing is a Multi-dimensional Problem

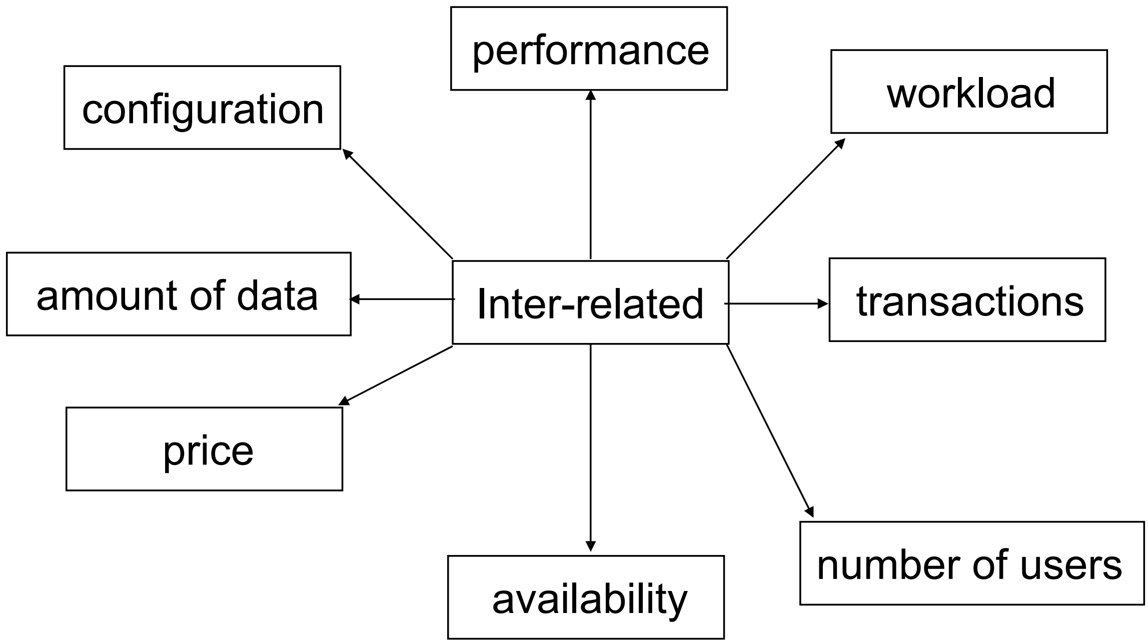

Figure 9-10. Sizing is a Multi-dimensional Problem AM711.0

Notes:

A complicating factor is that sizing is a multi-dimensional problem and the dimensions are interrelated. Consider these examples:

- Adding more hardware to address performance issues increases the cost.
- Adding more disks may require more adapters and perhaps more memory.
- Adding more memory can reduce the use of disks and save processor power.

In order to manage these inter-related sizing problems what is needed is a structured and practical approach.

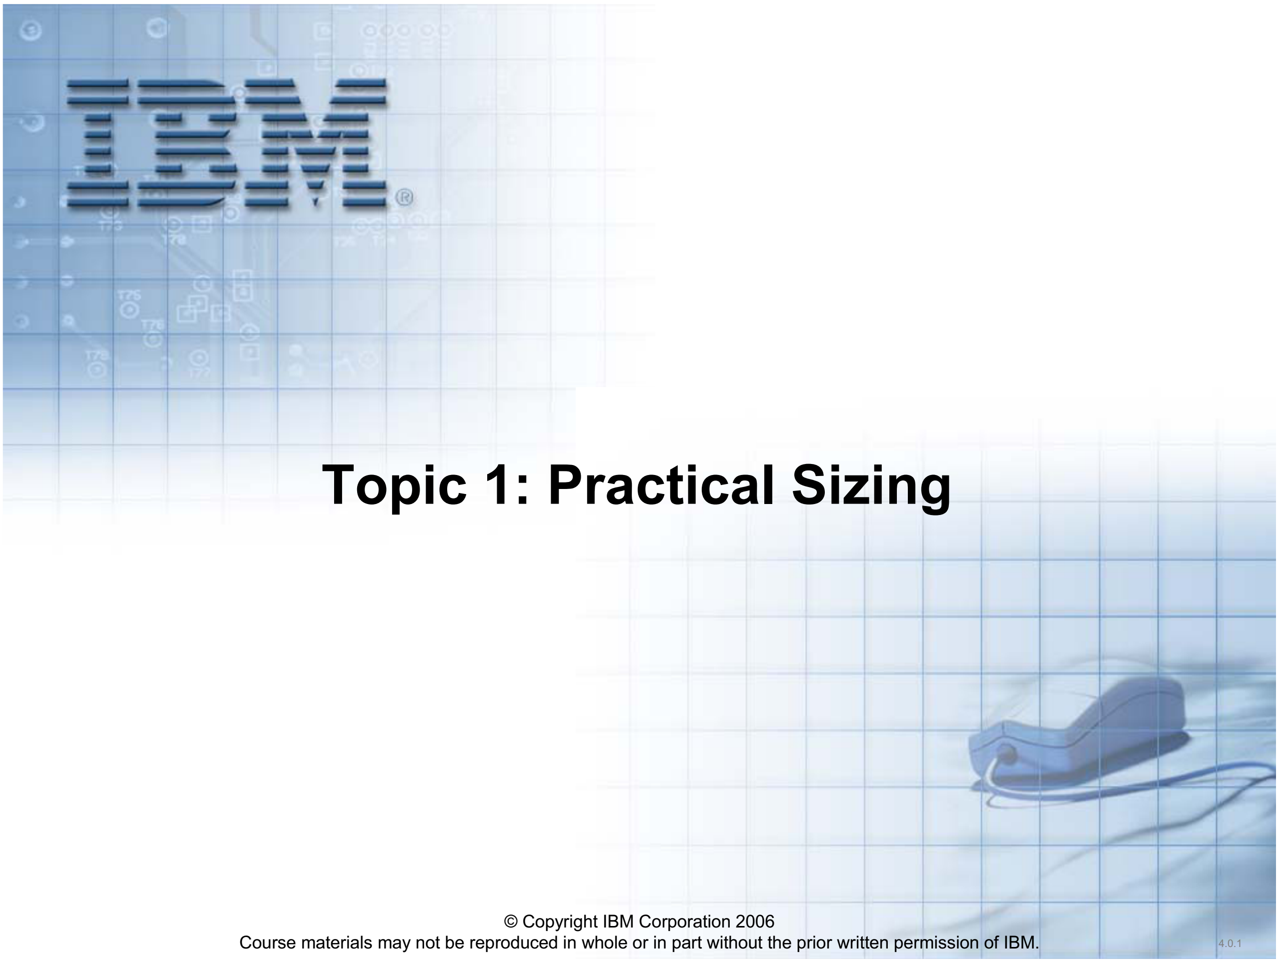

Topic 1: Practical Sizing

4.0.1

Figure 9-11. Topic 1: Practical Sizing AM711.0

Notes:

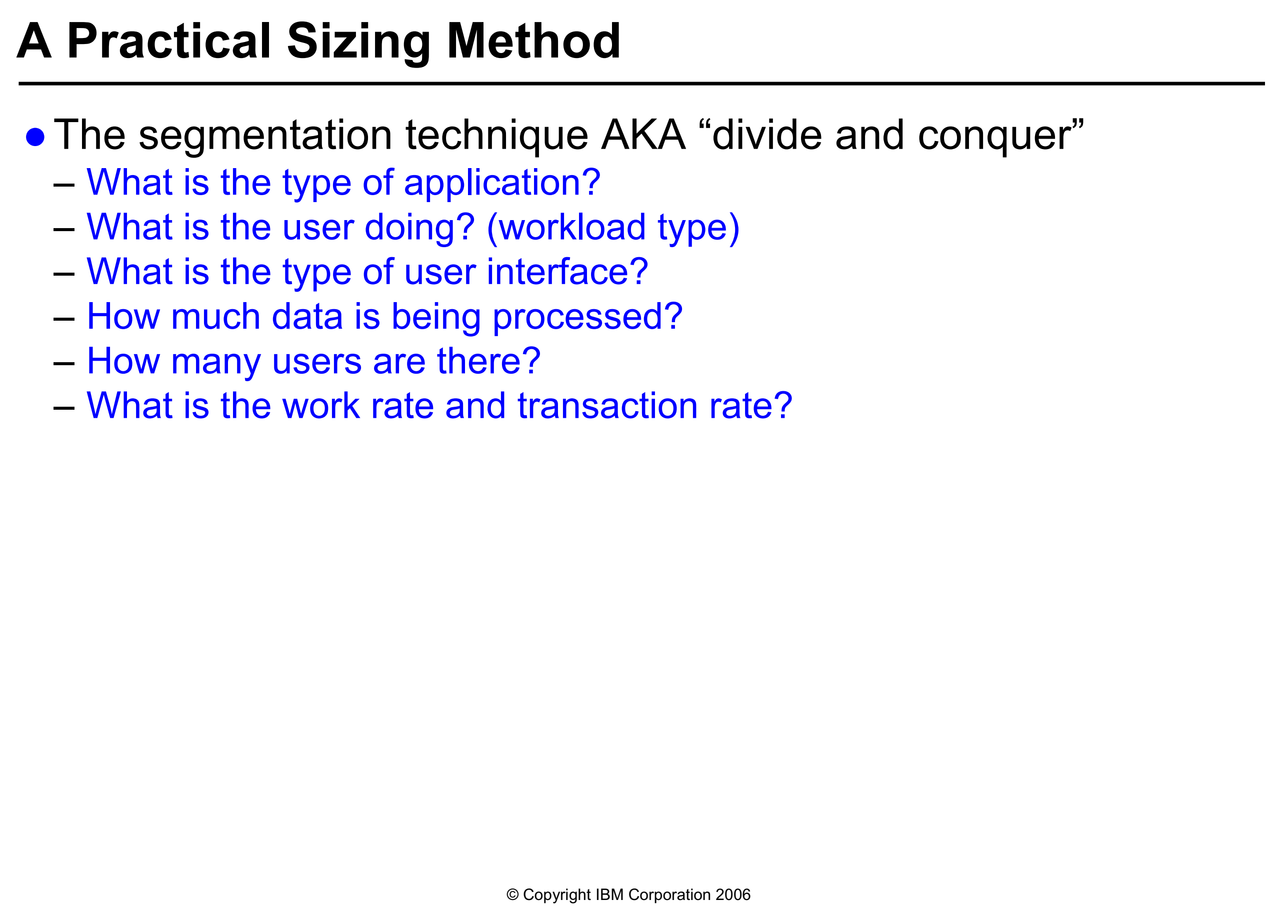

A Practical Sizing Method

- The segmentation technique AKA “divide and conquer”
  - What is the type of application?
  - What is the user doing? (workload type)
  - What is the type of user interface?
  - How much data is being processed?
  - How many users are there?
  - What is the work rate and transaction rate?

Figure 9-12. A Practical Sizing Method AM711.0

Notes:

This approach is useful for breaking down a sizing problem into manageable pieces. It involves asking questions that the client may actually be able to answer. Asking questions about the transaction size in terms of processor clock cycles or the per user memory requirement in bytes for code, stack and heap may not yield any information at all. The answers to the questions stated above will reduce the problem and assist in making a general guess. Let us look at the questions in more detail.

What is the Type of Application?

The application may be:

- RDBMS server
- Application server
- Stand-alone server (both of the above)
- Web server
- CAD/CAM workstation
- Multi-user system
- File/print server
- E-business
- Lotus Notes or similar workgroup application

Figure 9-13. What is the Type of Application? AM711.0

Notes:

The application type can help the person who is performing the sizing to understand what data is required and likely to be available.

What is the User Doing (Workload Type)?

The workload type may involve:

- Online transaction processing (OLTP)
- Interactive
- Enterprise Resource Planning (ERP) or Customer Relationship Management (CRM)
- Business Intelligence (BI) analysis
- Batch
- Web Access

Figure 9-14. What is the User Doing (Workload Type)? AM711.0

Notes:

This user or workload type can help the person who is performing the sizing since each has an associated size of transaction and likely numbers of users.

What is the Type of User Interface?

The type of user interface may be:

- Web browser
- PC-based application
- PC-based terminal emulation
- X Window System
- Dumb terminal
- Batch only

Figure 9-15. What is the Type of User Interface? AM711.0

Notes:

The interface type helps the person who is performing the sizing by indicating whether the user's application is running on the pSeries system or on local workstations or PCs.

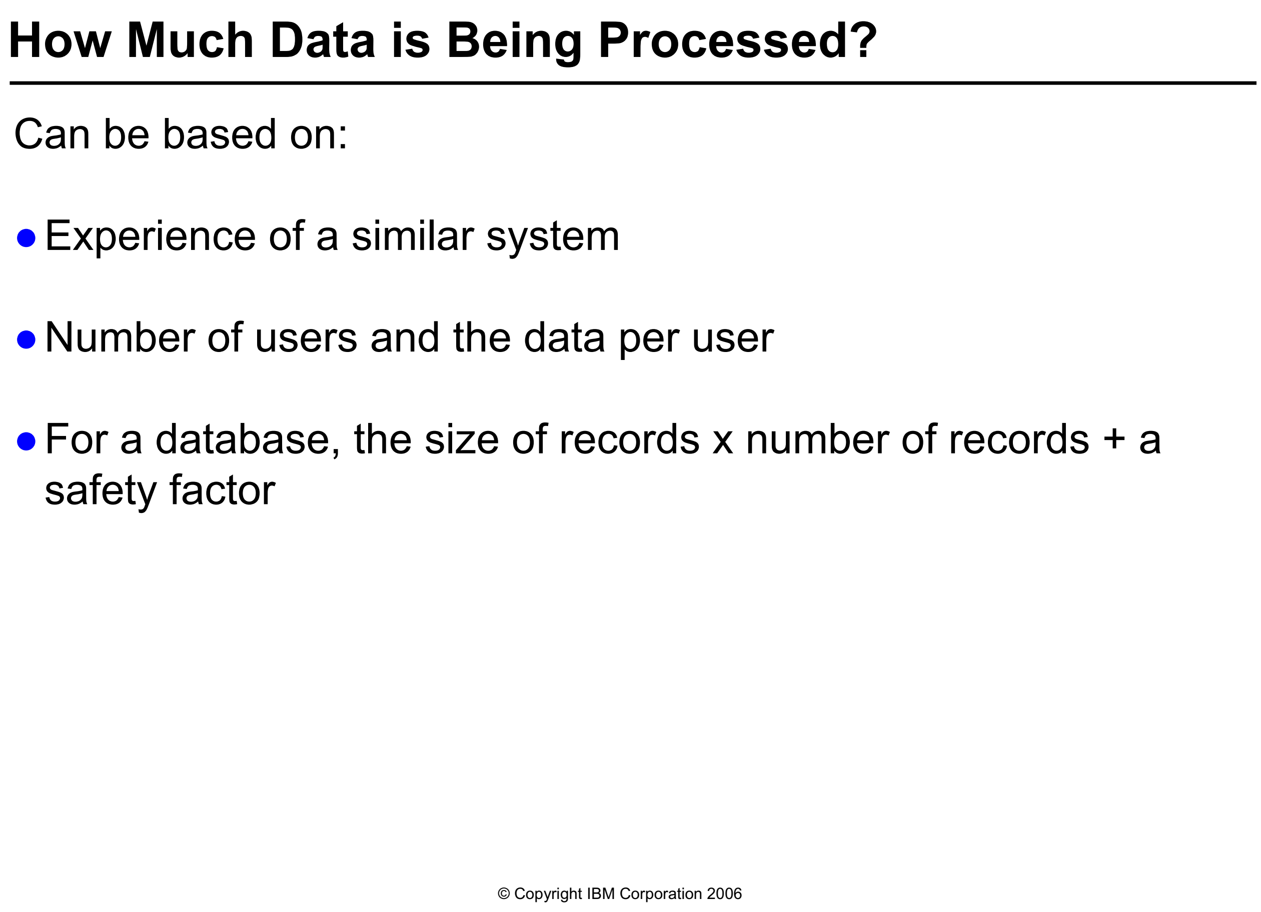

How Much Data is Being Processed?

Can be based on:

- Experience of a similar system
- Number of users and the data per user
- For a database, the size of records x number of records + a safety factor

Figure 9-16. How Much Data is Being Processed? AM711.0

Notes:

For the RDBMS establish whether the amount of data is the raw data or the disk size, and include the units (MB, GB, TB). For an RDBMS the ratio is typically on raw data to three disks.

What is the Number of Users?

There are many ways to count users.

- For sizing purposes only consider:
  - The number of users who actively interface with an application running on the system during the peak period.
  - The number of users who log on each day.

Figure 9-17. What is the Number of Users?

AM711.0

Notes:

Number of users may be expressed as:

- The number of people in the group although some of them never use the system.
- The number of people who have an account and can log on.
- The number of people who log on each day, but may not actually use it at the same time.
- The number of people who use the system at least once per day, but not regularly all day.

None of these methods of calculating the number of users is suited to the sizing process.

What is the Work Rate and Transaction Size?

- Difficult to generalize for UNIX servers:
  - OLTP = two seconds to a minute each
  - ERP applications = 2 to 10 minutes each
  - Web or e-business = zero seconds to one second
  - BI or DSS = 30 seconds to two weeks

- Transaction size and work rate is of prime importance in sizing because they determine the workload.
  Workload = TransactionSize x TransactionRate

Figure 9-18. What is the Work Rate and Transaction Size? AM711.0

Notes:

This is the last and most important fact to determine. Determining the work rate and transaction size is always a serious problem on UNIX-based systems. This is due to the general computing nature of the UNIX platform, the abundance of databases, variety of languages in which to write applications, different application types, and the range of transaction sizes and the user interfaces. A UNIX transaction is not a well quantified or understood concept. There is no such thing as a “normal” UNIX:

- Application
- Transaction
- Workload
- Query

There is no such thing as a user work rate either for:

OLTP = two seconds to a minute each

ERP applications = 2 to 10 minutes each

Web or e-business = zero seconds to as much as one second since it is a “hit” with no permanent concept of a user

BI or DSS = 30 seconds to two weeks

Processor requirements for these user types can easily vary by a factor of 1,000 (or 10,000 for BI applications). Also, high transaction rates often make higher disk input/output (I/O) demands too, so it affects more than just the processor requirements.

Most of the work in sizing involves attempting to characterize the processor workload because it is the item with least information. When it is not right, it causes the most painful upgrades. **Important:** Sizing mostly involves characterizing the processor workload.

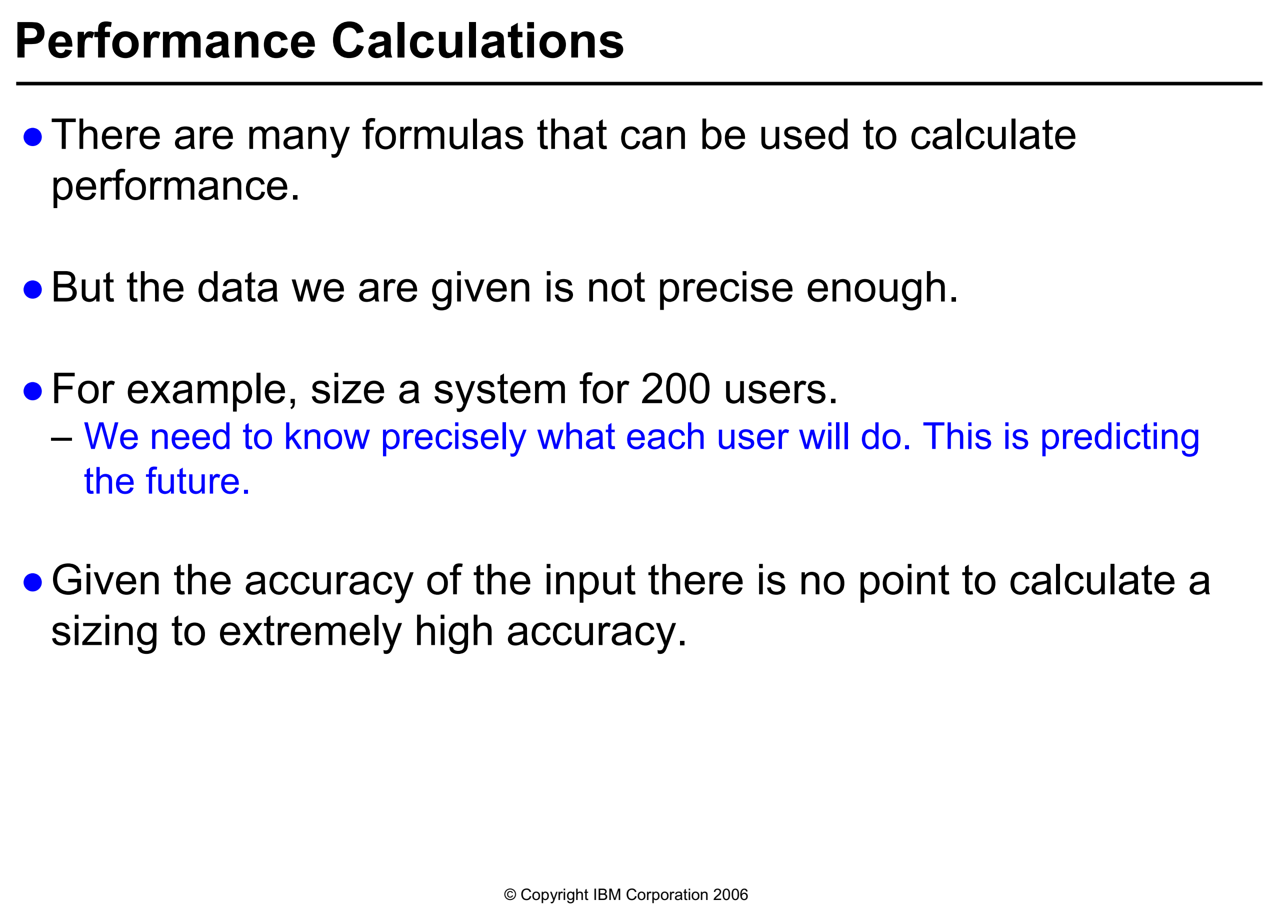

Performance Calculations

- There are many formulas that can be used to calculate performance.
- But the data we are given is not precise enough.
- For example, size a system for 200 users.
  - We need to know precisely what each user will do. This is predicting the future.
- Given the accuracy of the input there is no point to calculate a sizing to extremely high accuracy.

Figure 9-19. Performance Calculations AM711.0

Notes:

Even after we have segmented the sizing requirements there may still be too little information or the information may be too imprecise to use as a basis for calculating response times or the effect of changing workloads or hardware resources. Often the reasons for this are:

- The requirements are not accurately stated.
- The application resource use and the effects of usage on it, such as caching, have not been measured to sufficient accuracy.
- There is not sufficient time or no one is willing to fund it.

There are similar problems with resizing. Often the accuracy of information that is needed to use the performance calculations appropriately is not available because:

- What is happening on the system cannot be determined with high accuracy without drastically affecting the performance. This may be unacceptable to the client.
- Growth estimates are not accurate.

- There is not sufficient time or no one is willing to fund it.

Rules of Thumb

- Over the years sizing experts have built up a body of information about sizing from:
  - Experience
  - Benchmarks
  - Performance tuning
- This information can be used as ‘rules of thumb’
  - To estimate sizing when information is unavailable
- The goal is to size for balanced performance

Figure 9-20. Rules of Thumb AM711.0

Notes:

Sizing rules of thumb are based on the experience of sizing experts. They provide a basis on which to form sizing estimates which, while they may not be based on detailed information about the system, can be used practically and with a reasonable degree of accuracy.

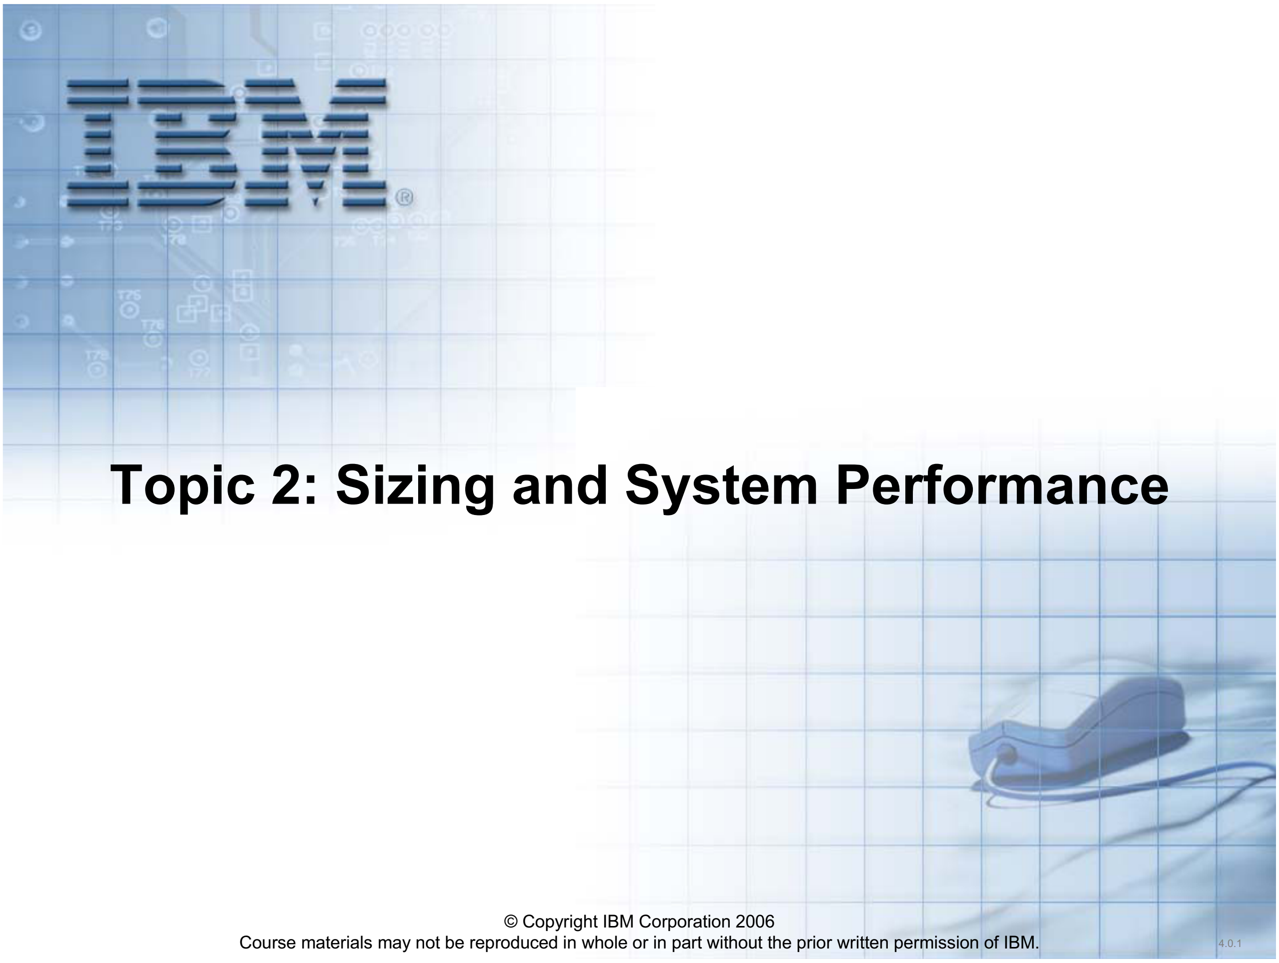

Topic 2: Sizing and System Performance

4.0.1

Figure 9-21. Topic 2: Sizing and System Performance AM711.0

Notes:

Balanced System Performance

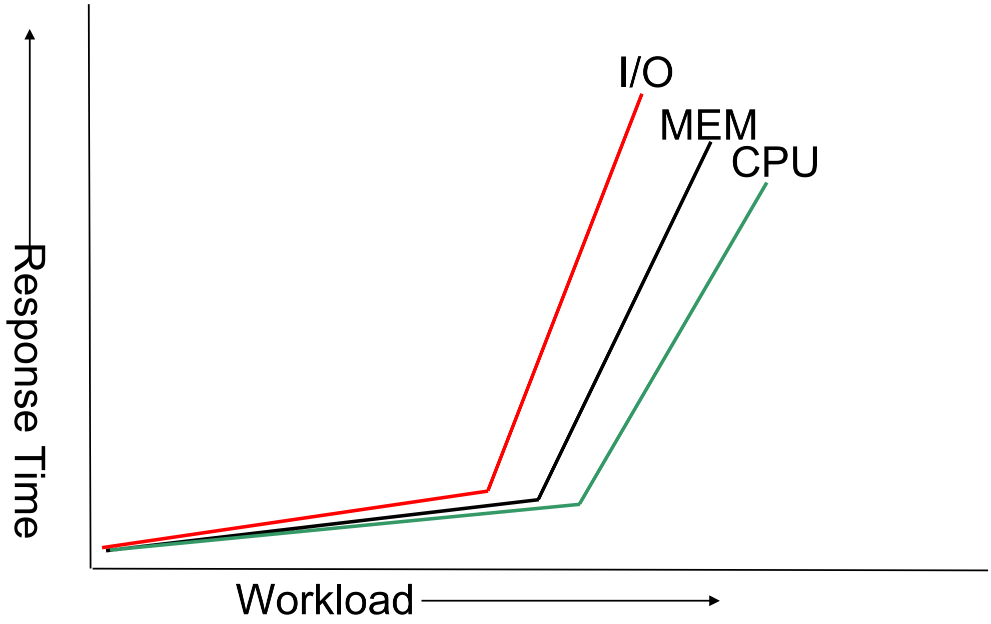

Figure 9-22. Balanced System Performance AM711.0

Notes:

In this simple diagram we see that the system response time increases, (a bad thing) as the workload increases. We also see that the subsystem components labelled as I/O, Memory, and CPU have similar curves, that is the performance of each subsystem gradually increases until it reaches some threshold and then performance (as measured by response time) degrades quickly. The diagram also shows that for this particular system the I/O performance reaches its threshold first. We would say that on this system the I/O is the bottleneck. It can be seen, however, that removing the I/O bottleneck, which could consist of adding more network or disk adapters or more spindles, or a variety of measures will eventually expose the next 'bottleneck' which is the memory subsystem. The very act of improving the I/O subsystem may in fact cause heavier memory utilization and expose the memory bottleneck earlier as the workload is increased. The aim is to provide a system in which the point at which each of the subsystem components reaches its threshold is as far along the workload axis as possible and as close to the other subsystems as possible. There is no advantage in having systems with incredible power at the CPU level but whose performance from an I/O perspective is marginal.

Utilization: The Good, the Bad, and the Ugly

- Rules of thumb exist for utilization of hardware components.
  - Good is the ideal value.
  - Bad is the point at which performance is likely to be affected and, if exceeded, the point at which performance is drastically affected.
  - Ugly is the point at which this component of system performance becomes a ‘hot spot’ or bottleneck.

Figure 9-23. Utilization: The Good, the Bad, and the Ugly AM711.0

Notes:

Utilization (1 of 2)

- Processor
  - **Good** 75% or less busy
  - **Bad** 85% busy
  - **Ugly** 90% or more busy

- Disk
  - **Good** 30% or less busy
  - **Bad** 40% busy
  - **Ugly** 50% or more busy

- Memory paging
  - **Good** zero paging
  - **Bad** 20 pages/sec/proc*
  - **Ugly** > 20 pages/sec/proc

- Network
  - **Good** 30% or less **
  - **Bad** 30% to 40% busy
  - **Ugly** 40% or more busy

Figure 9-24. Utilization (1 of 2) AM711.0

Notes:

* Large systems can page at this level and still perform very well. For example POWER4 systems can handle up to 20 pages/sec/proc.

** 30% or lower of the theoretical bandwidth. This level stops collisions on an Ethernet type network from becoming a problem.

Utilization (2 of 2)

- Run queue length
  - **Good** less than 2x the number of processors in length*
  - **Bad** 2x to 10x the number of processors in length
  - **Ugly** more than 10x the number of processors in length
- Paging space utilization**
  - **Good** less than 50% used
  - **Bad** 50% to 80% used
  - **Ugly** > 80%used***

Figure 9-25. Utilization (2 of 2) AM711.0

Notes:

* This allows each processor to finish its current task and immediately start another.

** The question is how big to make the paging space size? Unfortunately there is no one answer for this.

- For RDBMS systems use the greater value of 1GB or 25% of memory size.
- For systems with uncontrolled users (they can run whatever programs they like) use the greater value of 2GB or 100% of memory size.

*** This can get really ugly since running out of paging space can lead to processes being killed.

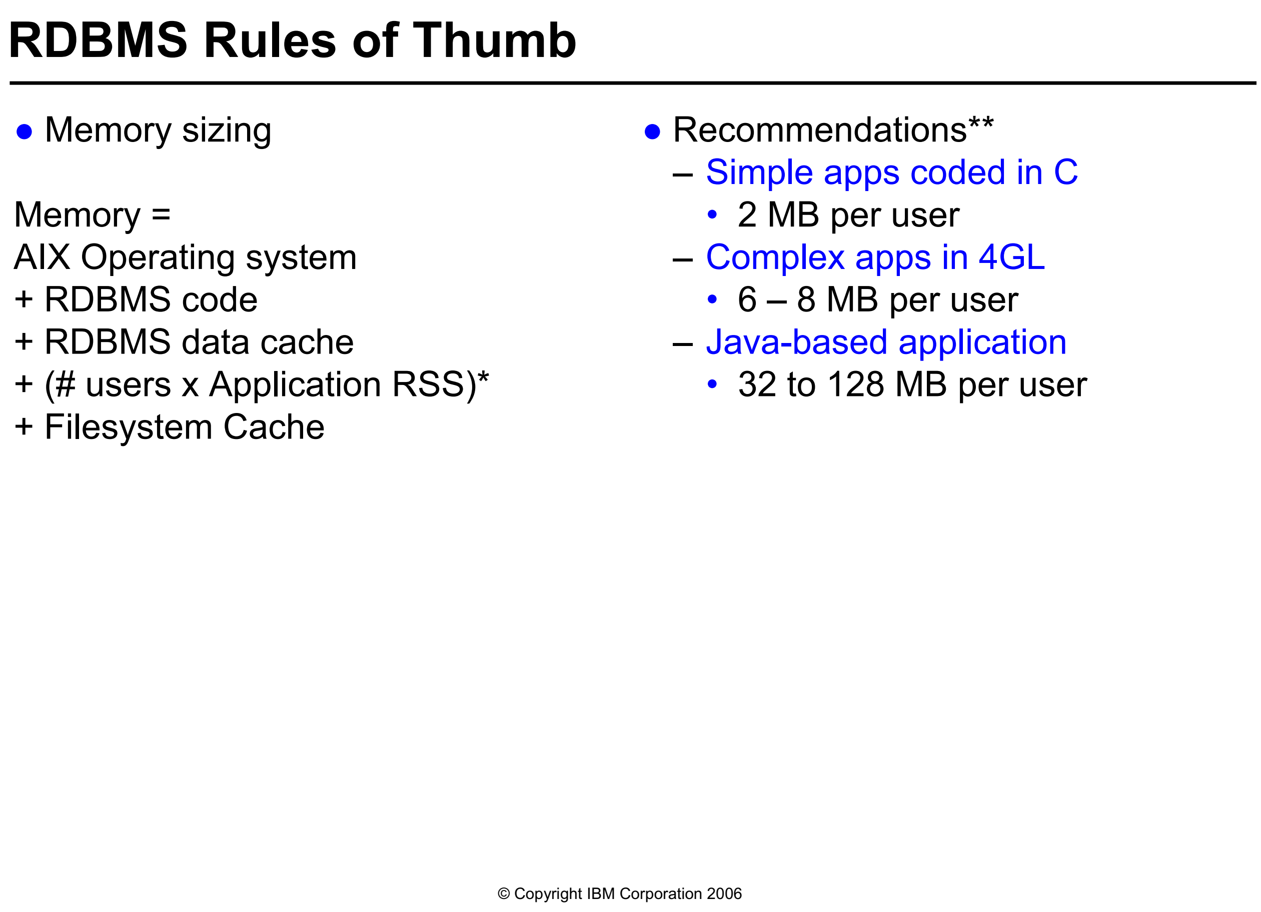

RDBMS Rules of Thumb

- Memory sizing

Memory =
AIX Operating system
+ RDBMS code
+ RDBMS data cache
+ (# users x Application RSS)*
+ Filesystem Cache

- Recommendations**
  - Simple apps coded in C
    - 2 MB per user
  - Complex apps in 4GL
    - 6 – 8 MB per user
  - Java-based application
    - 32 to 128 MB per user

Figure 9-26. RDBMS Rules of Thumb AM711.0

Notes:

The following sizing rules of thumb apply not only to RDBMS systems but to any system having a workload with components which will use CPU, memory and I/O. If the workload is more processor intensive, such as an application server, then ignore the disk I/O details.

*Application Resident Set Size is the code and data of the application that each user actually needs in memory to run. The resident set refers to the fact that operating systems such as AIX and Linux do not need to have the entire program in memory in order to run. Usually only a portion of the application is required. Also when AIX and Linux have processes running the same application, the code is shared and therefore saves memory.

** It would be preferable to measure the actual values

RDBMS Data and File System Cache (1 of 2)

- AIX uses a file system cache to speed up reading and writing files to and from file systems.
  - If RDBMS data is stored in file systems then a large file system cache is required.
  - If RDBMS data is stored in 'raw' logical volumes then the file system cache can be reduced in favor of more RDBMS cache.

Figure 9-27. RDBMS Data and File System Cache (1 of 2) AM711.0

Notes:

RDBMS Data and File System Cache (2 of 2)

- With no information use 50% RAM for RDBMS cache.
- In OLTP systems caching requirements are greater.
- May require 70% of RAM to ensure that tables and rows most often used are in memory.
- For OLTP as a starting point use 5% of data disk size for the cache
- For example, 500 GB raw data requires 25 GB of memory.
- BI and DSS use less memory and can work with only 30% of RAM for the RDBMS cache.

Figure 9-28. RDBMS Data and File System Cache (2 of 2) AM711.0

Notes:

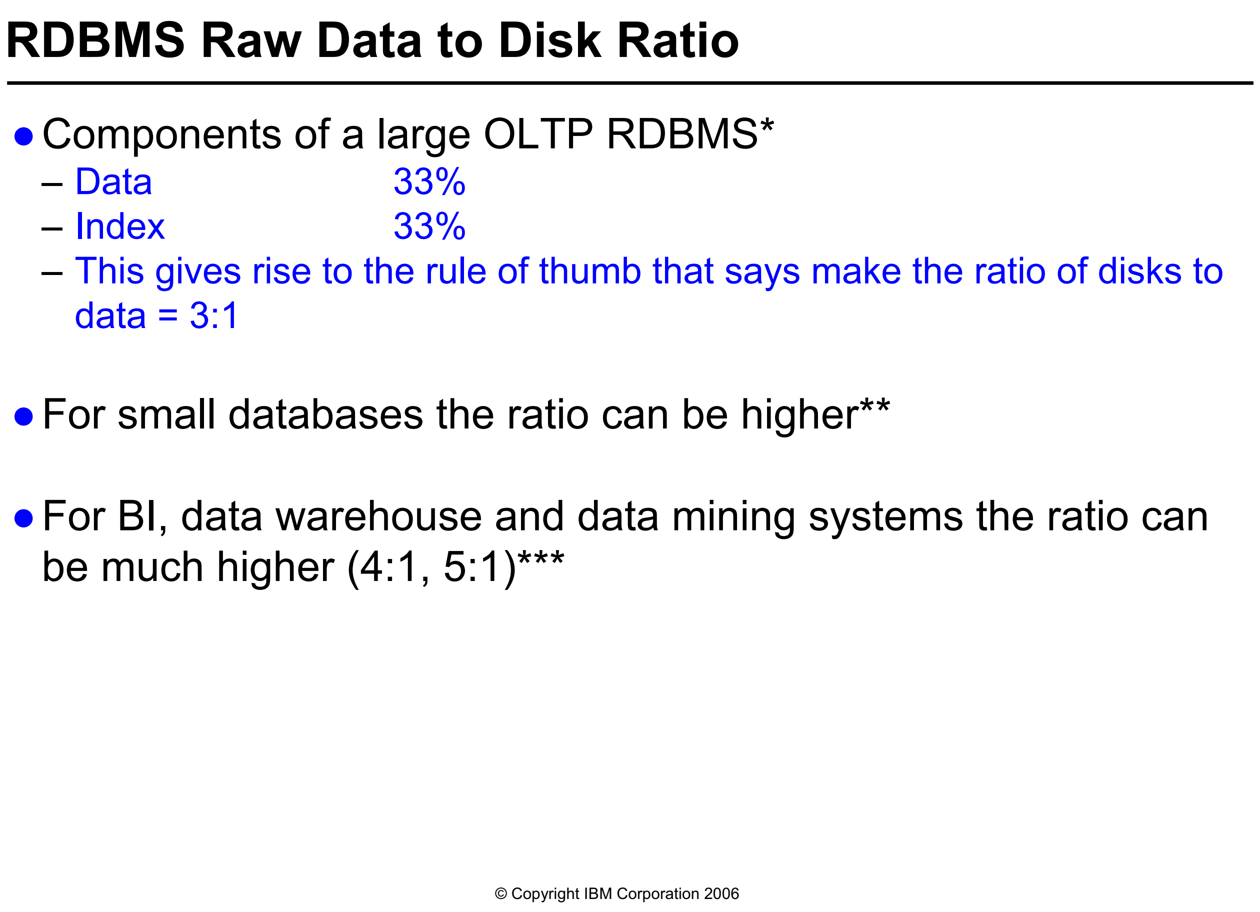

RDBMS Raw Data to Disk Ratio

- Components of a large OLTP RDBMS*
  - Data 33%
  - Index 33%
  - This gives rise to the rule of thumb that says make the ratio of disks to data = 3:1
- For small databases the ratio can be higher**
- For BI, data warehouse and data mining systems the ratio can be much higher (4:1, 5:1)***

Figure 9-29. RDBMS Raw Data to Disk Ratio AM711.0

Notes:

* For large databases on average only 1/3 of the size of the database is actually data. The remaining 2/3 is comprised equally of the index and temporary space. Note that this is an average and not a rule. If you use the 3:1 data to disk ratio it must be documented as an assumption.

** The small database needs to be configured with enough disks to be safe, for example, RAID 7+1.

*** BI environments need lots of disk space for the Extraction, Transform, Load (ETL).

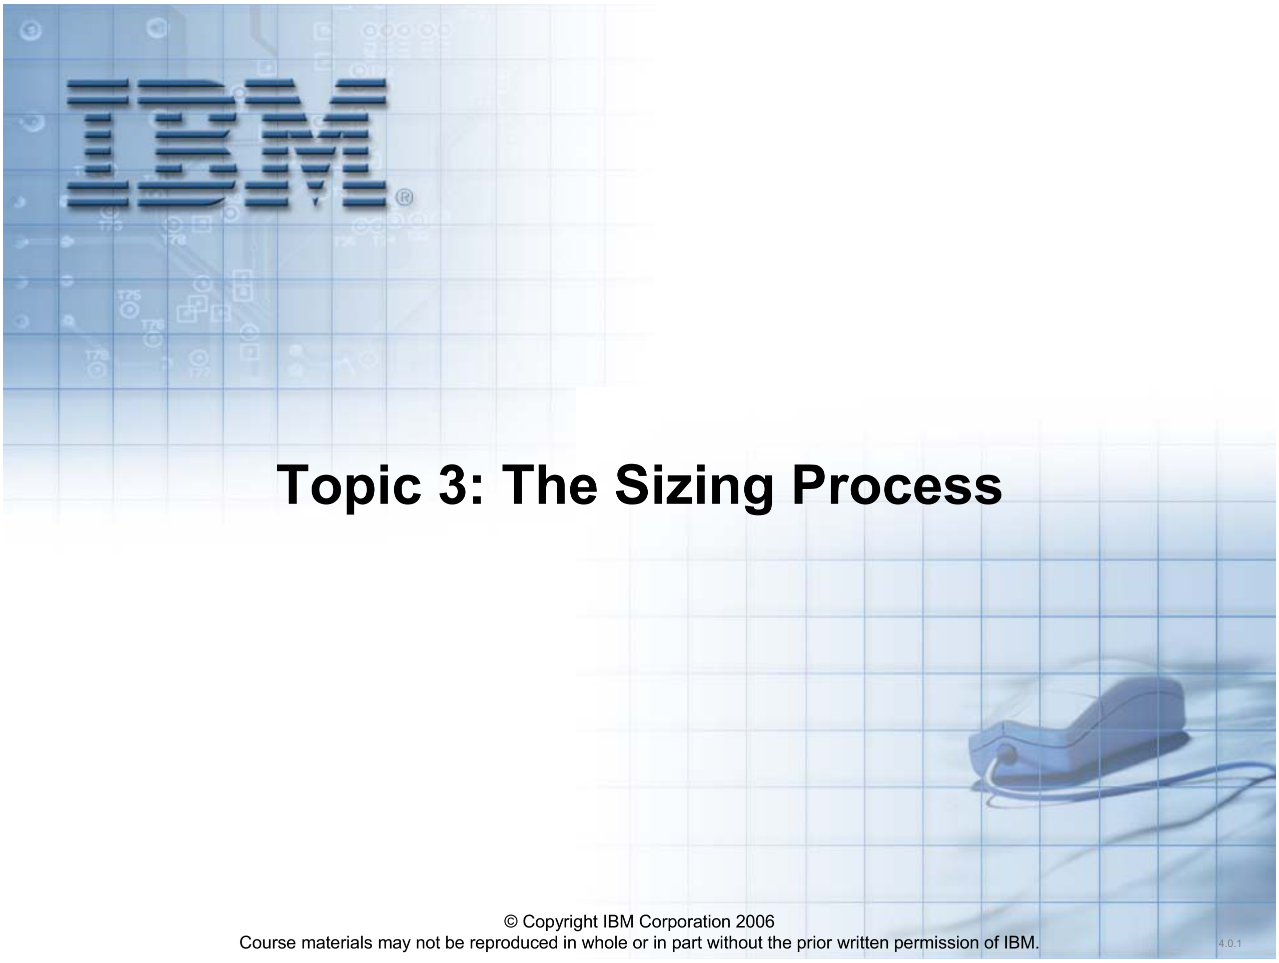

Topic 3: The Sizing Process

4.0.1

Figure 9-30. Topic 3: The Sizing Process AM711.0

Notes:

Sizing and Assumptions

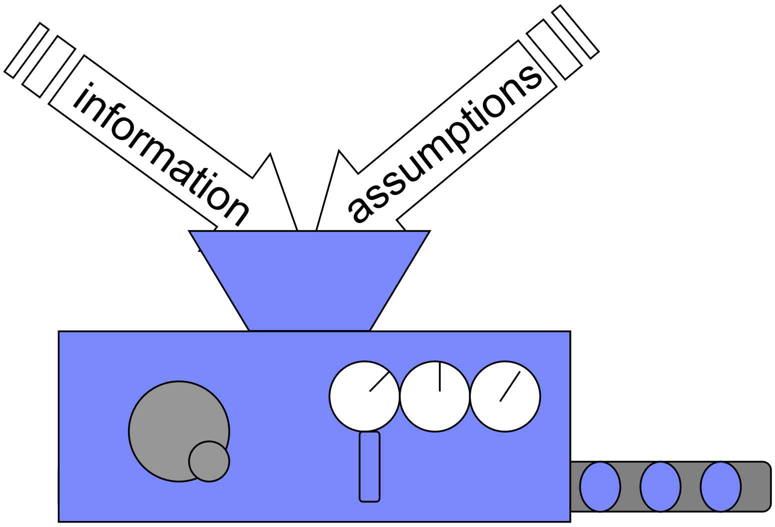

Figure 9-31. Sizing and Assumptions AM711.0

Notes:

It isn't necessarily the assumptions that cause the problems it is the implicit, hidden assumptions.

Sizing Tasks

- Solution Design and Requirements
- Sizing Model
- Hardware Requirements
- Building Block Choices
- eConfig/Pricing

Figure 9-32. Sizing Tasks AM711.0

Notes:

Sizing Inputs

- Application top-level design
- Raw data or disk size in GB
- Numbers of records and record sizes
- User numbers and user types
- Network details such as topology or throughput requirements
- Transaction type and size
- Growth in terms of company expectations, data size or number of users
- Disk protection and disk type preferences
- Preference from client for large SMP environment or cluster of smaller systems
- Preference for price versus availability

And lots more…

Figure 9-33. Sizing Inputs AM711.0

Notes:

The inputs listed above are not exhaustive but do reflect the range of sizing inputs that may be provided. The range is not just in terms of level of detail but also includes information such as the client's preference for a particular type of solution.

Sizing Outputs

- pSeries model with number of processors and speed
- Memory configuration
- Amount and type of disk drives
- Adapters for disk, tape and network
- Software requirements (operating system, applications, and so forth)
- Pricing and budget constraints

- Client project name, and contacts
- Documented input
- Documented assumptions
- Risk analysis
- Additional considerations
- Extras included

Figure 9-34. Sizing Outputs AM711.0

Notes:

Sizing output is generated during the sizing process. The list on the left are items which you should end up with in order to configure a system. The list on the right are other items which are also important for a successful design.

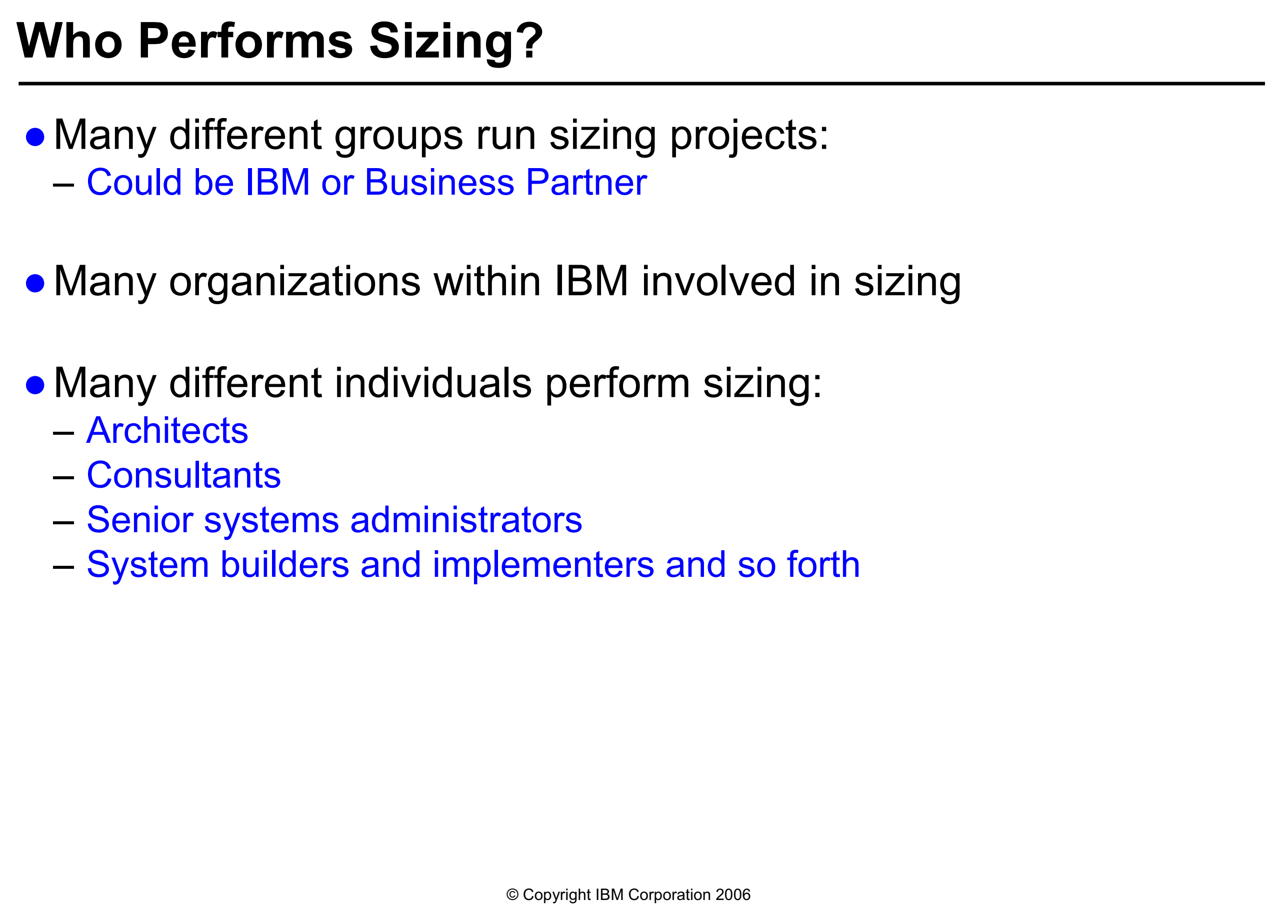

Who Performs Sizing?

- Many different groups run sizing projects:
  - Could be IBM or Business Partner
- Many organizations within IBM involved in sizing
- Many different individuals perform sizing:
  - Architects
  - Consultants
  - Senior systems administrators
  - System builders and implementers and so forth

© Copyright IBM Corporation 2006

Figure 9-35. Who Performs Sizing? AM711.0

Notes:

Who performs sizing

The groups who tend to run sizing projects are:*

IBM Server Group

- pSeries Sales
- pSeries Sales Technical Specialists
- IBM Techline Support
- Field Technical Sales Support (FTSS)
- Advanced Technical Support (ATS)

Clients, integrators, or IBM Global Services

- Architects
- Consultants
- Senior systems administrators

- System builders and implementers

*(Not an exhaustive list)

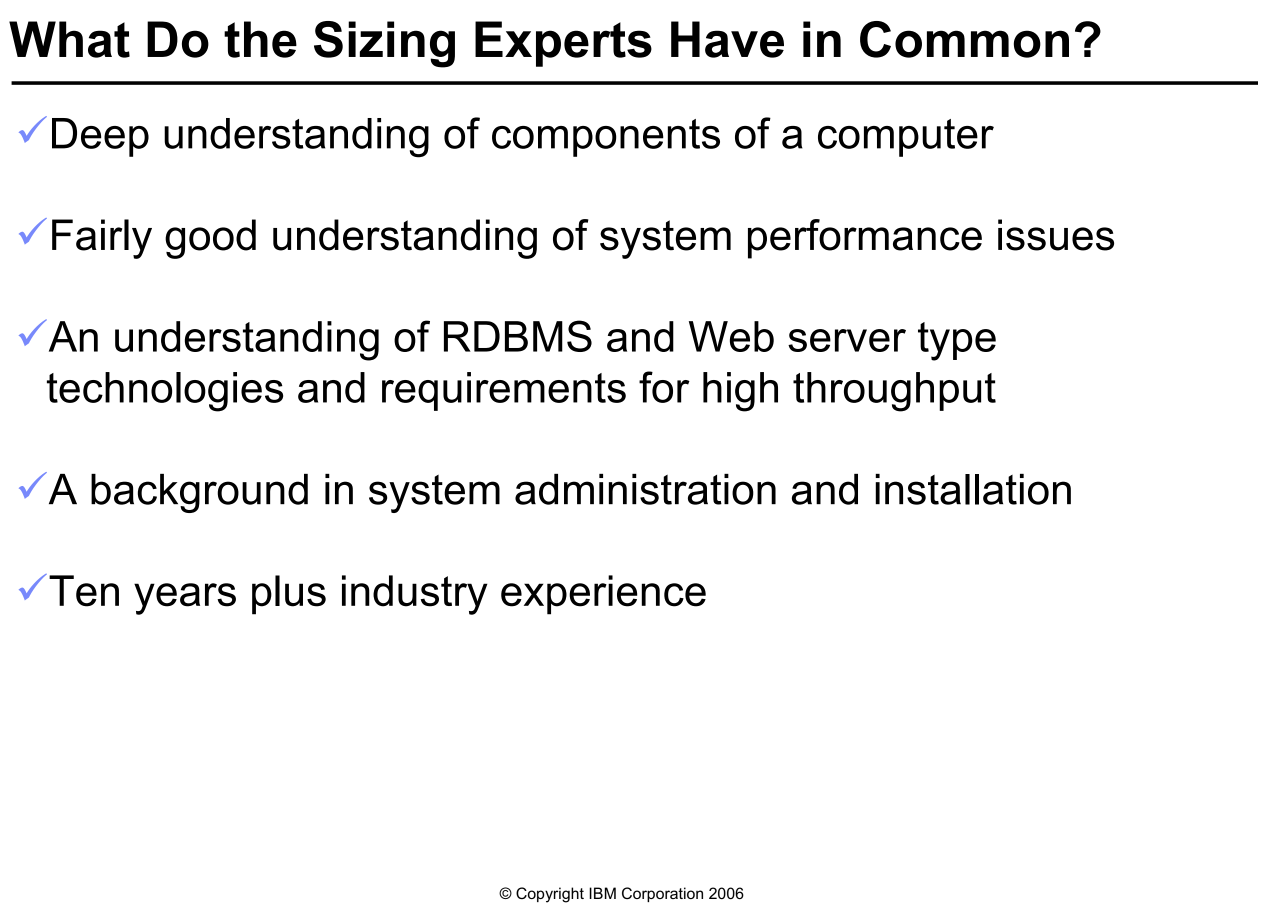

What Do the Sizing Experts Have in Common?

- ✓Deep understanding of components of a computer
- ✓Fairly good understanding of system performance issues
- ✓An understanding of RDBMS and Web server type technologies and requirements for high throughput
- ✓A background in system administration and installation
- ✓Ten years plus industry experience

Figure 9-36. What Do the Sizing Experts Have in Common? AM711.0

Notes:

These people all have the following characteristics in common:

A fairly deep understanding of the components of a computer and how they interoperate.

A fairly good understanding of system performance issues.

An understanding of RDBMS and Web server type technologies and what is required for high throughput.

A background in system administration and installation so they understand the issues of the groups that have to get the resulting computer system working.

At least ten years plus industry experience.

Sizing and Resizing Process – Sizing Tasks

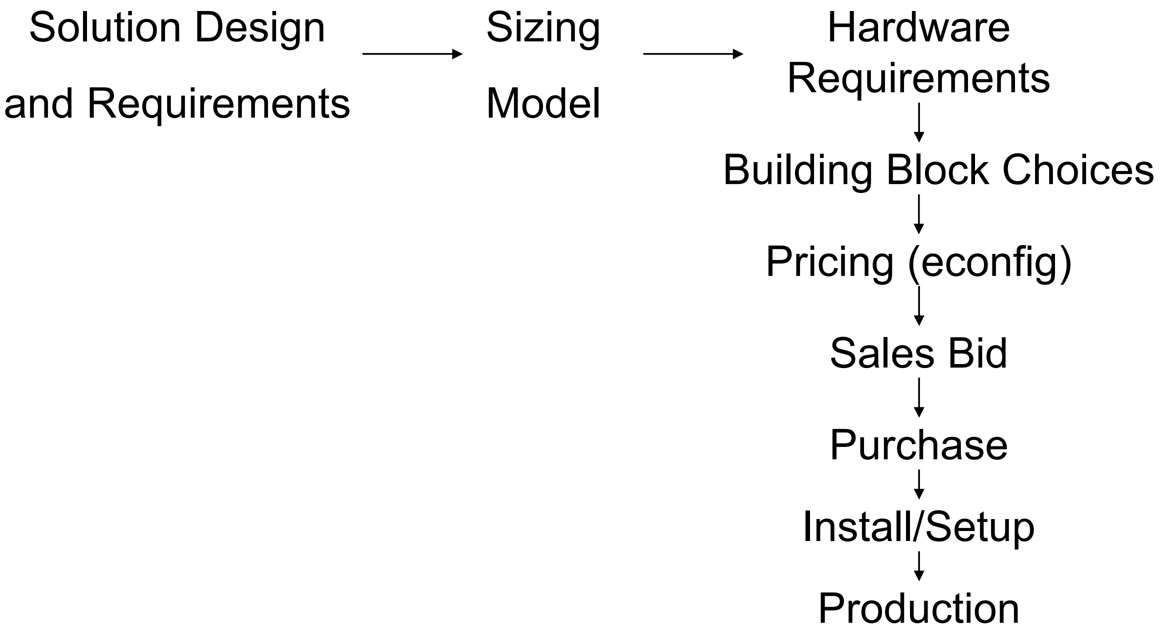

Solution Design and Requirements → Sizing Model → Hardware Requirements

Hardware Requirements
↓
Building Block Choices
↓
Pricing (econfig)
↓
Sales Bid
↓
Purchase
↓
Install/Setup
↓
Production

Figure 9-37. Sizing and Resizing Process – Sizing Tasks AM711.0

Notes:

Sizing and Resizing Process – Tuning Tasks

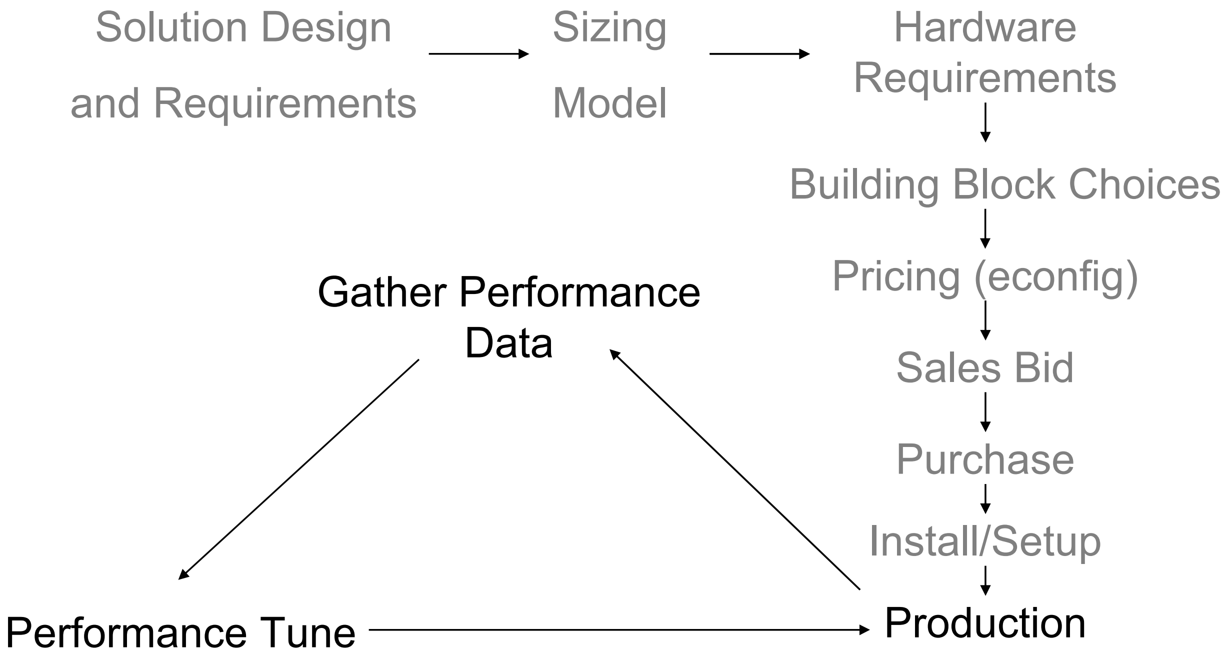

Solution Design and Requirements

Sizing Model

Hardware Requirements

Building Block Choices

Pricing (econfig)

Sales Bid

Purchase

Install/Setup

Production

Gather Performance Data

Performance Tune

Figure 9-38. Sizing and Resizing Process – Tuning Tasks AM711.0

Notes:

Sizing and Resizing Process – Resizing Tasks

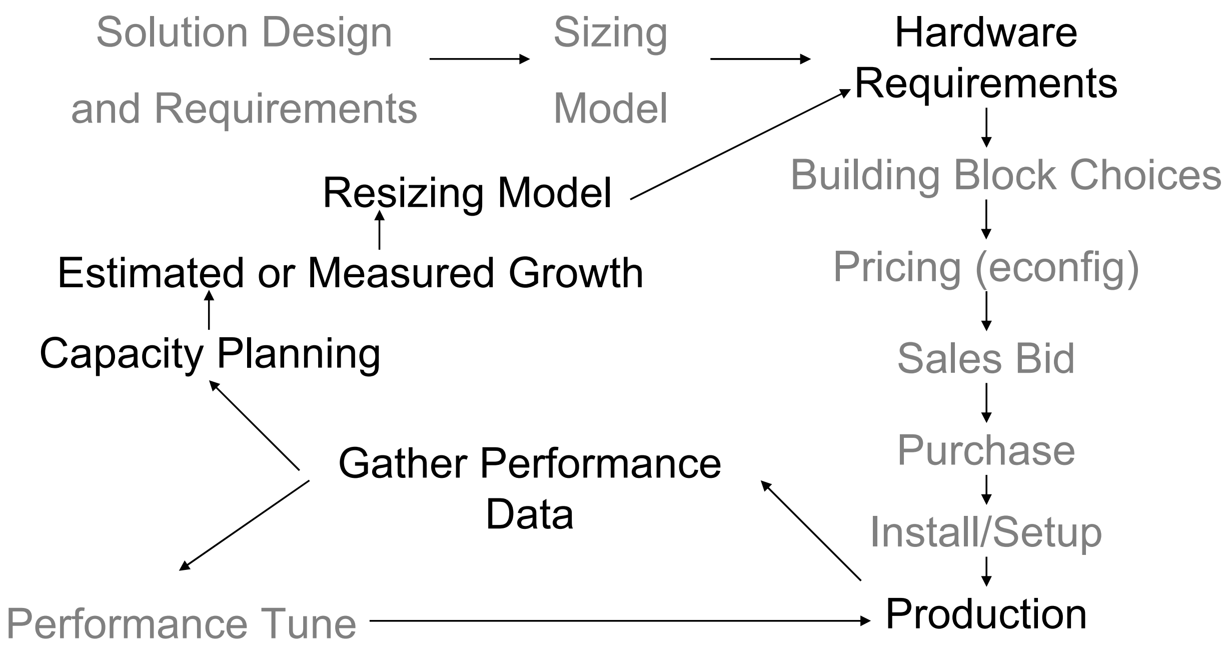

Solution Design and Requirements

Sizing Model

Hardware Requirements

Building Block Choices

Resizing Model

Estimated or Measured Growth

Pricing (econfig)

Capacity Planning

Sales Bid

Gather Performance Data

Purchase

Install/Setup

Performance Tune

Production

Figure 9-39. Sizing and Resizing Process – Resizing Tasks AM711.0

Notes:

Components Involved in Sizing

- Hardware Components
  - Processors
  - Memory
  - Input/Output
  - Storage
- Software Components
  - Operating System
    - AIX
    - Linux
  - Applications

Figure 9-40. Components Involved in Sizing AM711.0

Notes:

Hardware Requirements

- For each workload specify:
  - Processor power rating
  - Memory speed and size
  - Disk drives number and type
  - If scale out is possible/desirable

Figure 9-41. Hardware Requirements AM711.0

Notes:

Hardware Requirements

For each of the workloads, the hardware requirement is specified as:

Processor power rating (for example, estimated rPerf)

Memory: Speed and size

Disk Drives: Number and size

Adapters: Number and type

The Scale out column highlights when a workload can be spread between multiple, smaller servers. This is important for reducing costs and increasing availability.

Hardware Requirements - Example

Service name	CPU rating	RAM (GB)	Disks (36 GB)	Network (GB)	FC	Scale out
Database	50	64	200	2 + 1	8	
Application	80	120	1	4 + 1	2	yes
Web server	20	16	1	4 + 1		yes
Test server	40	60	80	3	4	
Development	15	32	20	1		
Backup	4	8	8	2	2	

Figure 9-42. Hardware Requirements - Example AM711.0

Notes:

The Scale out column highlights when a workload can be spread between multiple, smaller servers. This is important for reducing costs and increasing availability.

Building Block Choices

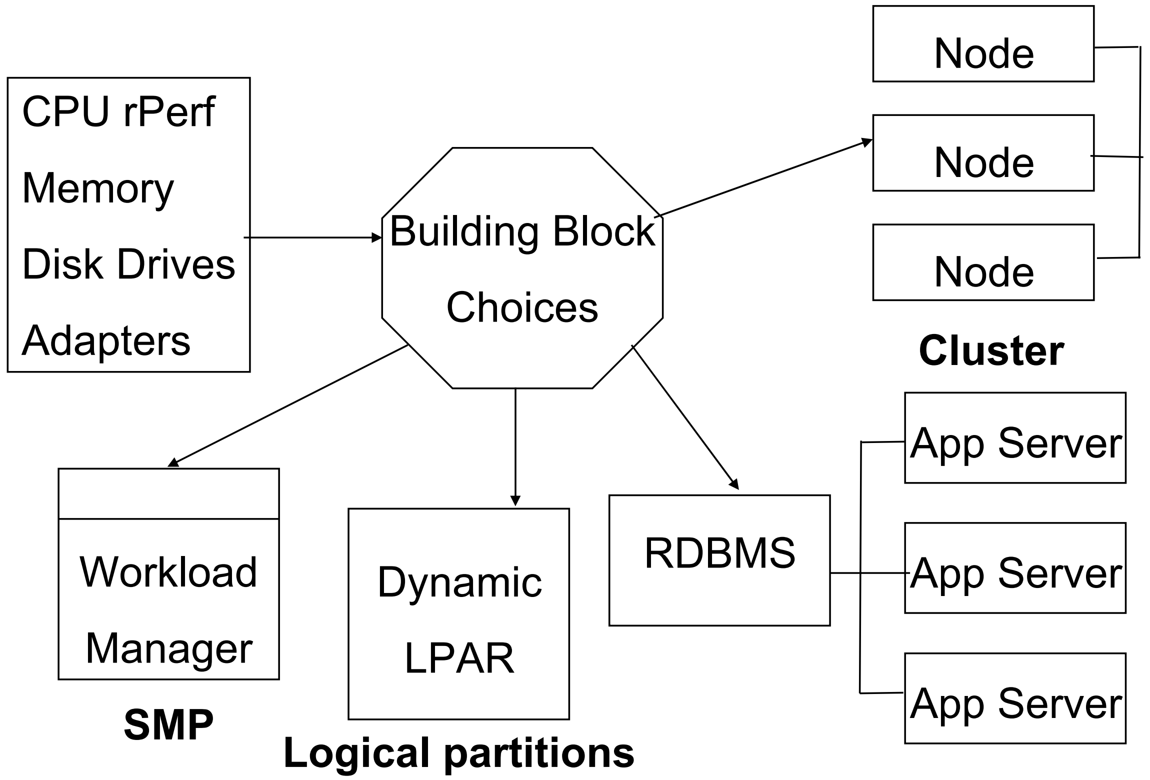

© Copyright IBM Corporation 2006

Figure 9-43. Building Block Choices AM711.0

Notes:

After the sizer determines the hardware requirements, you must select the systems to provide the resources. If this solution is tiered or has several independent components, then there is a hardware requirement for each. The next phase is to consider the building blocks that make up the system. Fortunately, the pSeries family of workstations and servers is extremely flexible. However, this flexibility means that someone has to make a decision about which one to choose. Often it is a given by the sizing request. For example, the request may state that the client wants either a single large system with logical partitions (LPAR) for maximum flexibility, or for cost reasons and the application's natural fit, they want a cluster of smaller systems. If there is no more guidance, the person who is performing the sizing must decide for themselves, discuss this further with the requester, or provide multiple solutions and options for the client to decide.

Weighing Sizing Components

- Concentrate on most expensive first
  - pSeries model
  - Memory configuration
  - Disk drive type and number
  - Adapters for disk, tape, and network
  - Software requirements
- Look to future requirements
  - Avoid the box swap

Figure 9-44. Weighing Sizing Components AM711.0

Notes:

Certain sizing components are more important, and in fact more critical, than others. If you calculate them incorrectly, they can be the quite painful to correct. The following components are important when sizing:

- pSeries model including number of processors and their speed
- Memory configuration
- Disk drive or drives' type and number
- Adapters for disk drives, tapes, and networks
- Software requirements

When trying to determine the optimal system size, concentrate on the most expensive component first. Make sure that, if necessary, there is a simple way to easily correct these components later and avoid replacing the entire system. Technical experts who perform sizing all agree that it is the pSeries model that is the most important component, and not just configuring the maximum number of the fastest processors. This means, if necessary, you can add a few more processors or faster processors later. Upgrading to the next model up (by way of a box swap) is expensive and time consuming for clients. It is best to avoid this.

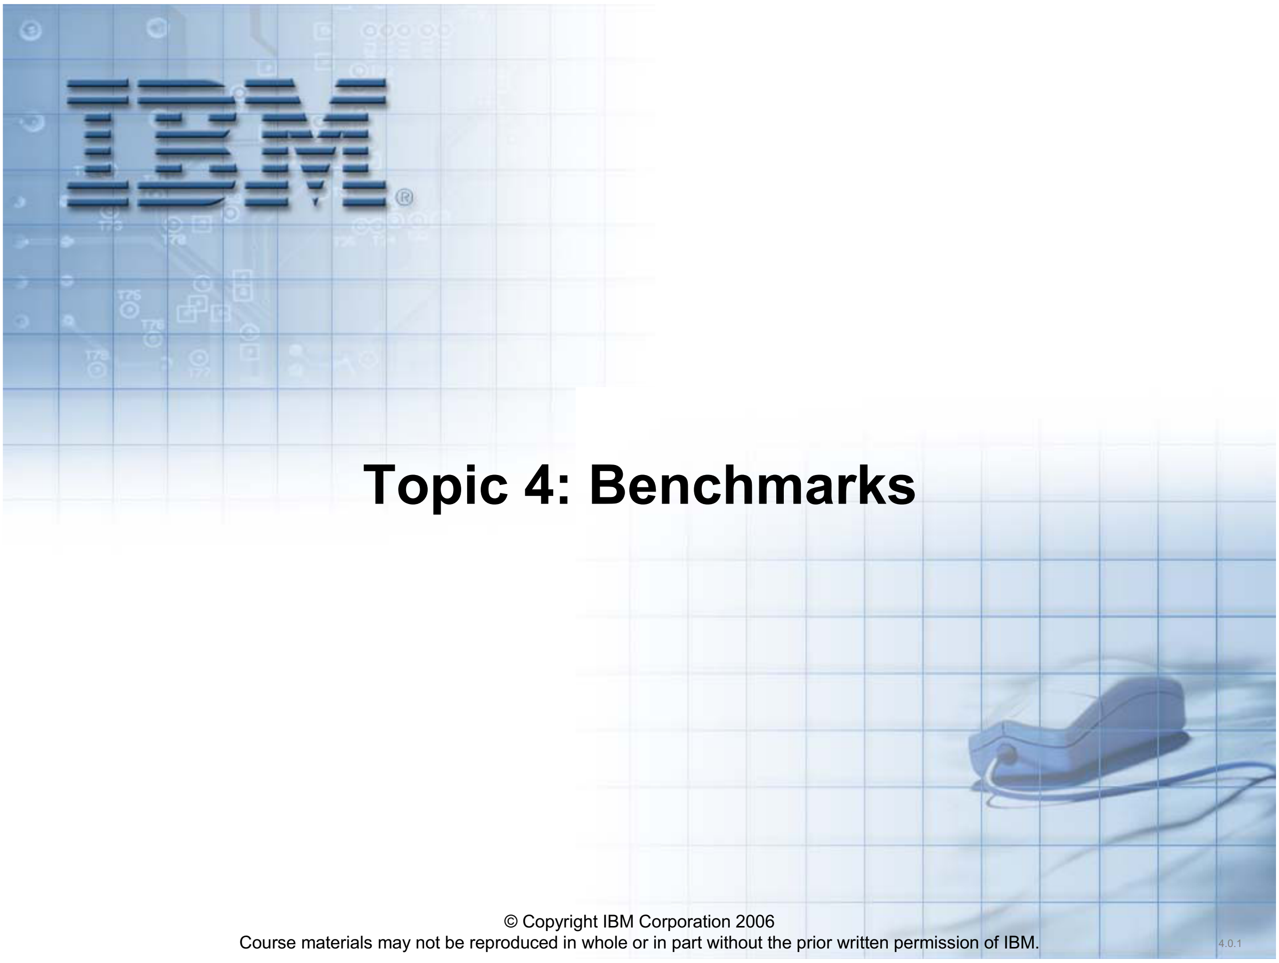

Topic 4: Benchmarks

Figure 9-45. Topic 4: Benchmarks AM711.0

Notes:

Benchmarks

- Online transaction processing (OLTP) benchmarks
- Business Intelligence (BI) benchmarks
- E-business benchmarks
- High Performance Computing (HPC) benchmarks
- Independent software vendor (ISV) benchmarks

Figure 9-46. Benchmarks

AM711.0

Notes:

Benchmarks and Solution Sizing

- Why study benchmarks?
- How are benchmarks different than ‘real’ performance?
- What conclusions may be drawn from benchmark results?

Figure 9-47. Benchmarks and Solution Sizing AM711.0

Notes:

There are several reasons to study benchmark tests. The most well-known reason is to compare performance of different computers because clients need to know which system is better for their applications when they are deciding to buy a system. However, it is difficult to find an absolute measurement because, nowadays, computers are complex systems in which many components influence the overall performance of the system. System performance especially depends on the kind of application software that is running on the system. Benchmarks are necessarily abstract and simplified models of application environments. For this reason, benchmarks represent a good measuring tool to compare different systems rather than a precise tool for capacity planning for a given client application environment. No benchmark can fully characterize the performance of a system in a true production environment because:

- The behavior of benchmark applications is essentially constant on a given system. Real applications, when executed several times, almost invariably have different inputs, and consequently exhibit different behavior each time.

- Benchmarks are executed under ideal circumstances. The benchmark is typically the only application that is executed on a system dedicated to a single purpose. For this reason, system overheads, such as paging and context switches, are lower than in actual production use of a processor.

Benchmark processors are often equipped with the latest and greatest memory and disk subsystems. They may include features that may not match exactly a system that is of interest to a client. In this sense, benchmarks represent the upper limit of system performance.

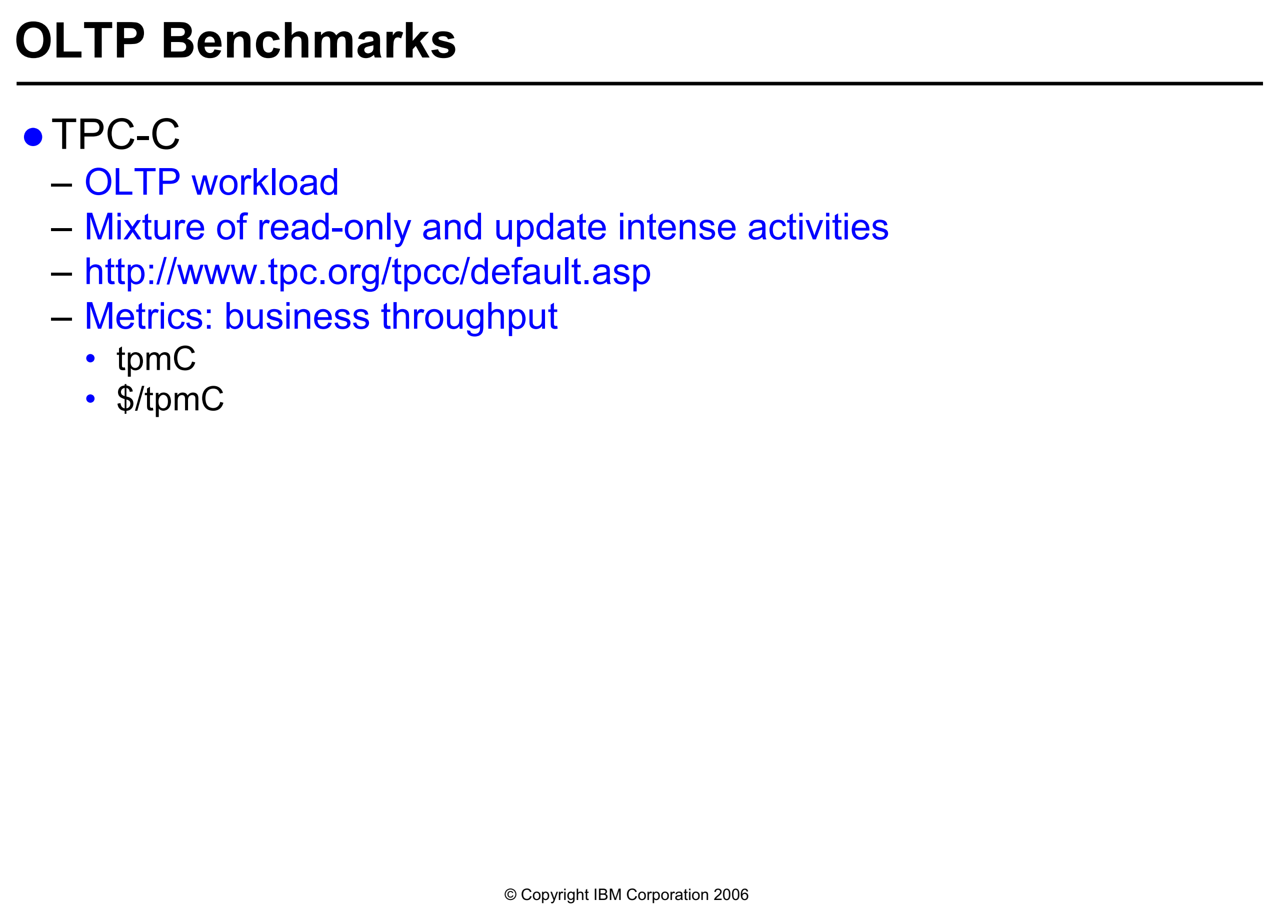

OLTP Benchmarks

- TPC-C
  - – OLTP workload
  - – Mixture of read-only and update intense activities
  - – http://www.tpc.org/tpcc/default.asp
  - – Metrics: business throughput
    - tpmC
    - $/tpmC

Figure 9-48. OLTP Benchmarks AM711.0

Notes:

TPC benchmark C (TPC-C) is an OLTP workload. It is a mixture of read-only and update-intensive transactions that simulate the activities found in complex OLTP application environments. It does so by exercising a breadth of system components associated with such environments, which are characterized by:

The simultaneous execution of multiple transaction types that span a breadth of complexity.

- Online and deferred transaction execution modes
- Multiple online terminal sessions
- Moderate system and application execution time
- Significant disk input/output (I/O)
- Transaction integrity: Atomicity, Consistency, Isolation, Durability (ACID) properties
- Nonuniform distribution of data access through primary and secondary keys
- Databases consisting of many tables with a wide variety of sizes, attributes, and relationships

- Contention on data access and update

Business Intelligence Benchmarks

- TPC-H
  - Ad-hoc decision support benchmark
    - Random queries
    - No cache benefit
  - Throughput and power tests
  - Metrics:
    - TPC-H Composite Query-per-hour Performance Metric (QphH@Size)
    - $/QphH@Size

Figure 9-49. Business Intelligence Benchmarks AM711.0

Notes:

TPC-H is an ad-hoc decision support benchmark that represents decision support environments. It is ad-hoc in the sense that queries are random.

Therefore no caching benefits influence the benchmark results. Given this, the query times can be quite long since you cannot tune the database for the query.

The benchmark consists of a suite of business-related ad-hoc and concurrent data modifications. TPC-H evaluates the performance of a decision support system that performs complex queries (more complex than OLTP transactions) on large volumes of data with a high degree of complexity.

TPC-H is composed of power and throughput runs. They should be executed under the same conditions:

- Power test: Measures the raw query execution power of the system when connected with a single active user.
- Throughput test: Measures the ability of the system to process the most queries in the least amount of time.

Systems today are used for both scale-up (supporting more users and higher throughput) and speed-up (making a single task faster, reducing response time) of a workload. The power metric demonstrates the speed-up while the throughput metrics show the scale-up capacity of the system.

The queries and the data populating the database were chosen to have broad industry-wide relevance, while maintaining a sufficient degree of ease of implementation. This benchmark illustrates decision support systems that:

- Examine large volumes of data.
- Execute queries with a high degree of complexity.
- Give answers to critical business questions.

e-business Benchmarks

- TPC-W benchmark
  - Transactional Web benchmark
  - Controlled Internet commerce environment
- Metrics:
  - Web interactions per second (WIPS)
  - System cost per WIPS ($/WIPS)
  - http://www.tpc.org/tpcw/default.asp

Figure 9-50. e-business Benchmarks

Notes:

Performed in a controlled Internet commerce environment that simulates the activities of a business-oriented transactional Web server. The workload covers a range of system components associated with such environments, which are characterized by:

- Multiple online browser sessions
- Dynamic page generation with database access and update
- Consistent Web objects
- The simultaneous execution of multiple transaction types that span a breadth of complexity
- Online transaction execution modes
- Databases consisting of tables with a variety of sizes, attributes, and relationships
- Transaction integrity (ACID properties)
- Contention on data access and update

SPEC JBB2000 Benchmark

- Measures Java server performance for business applications
  - Emulation of three-tier system
  - Business logic and object manipulation
- Metrics
  - Business operations per second (Ops/s) as an average number for peak performance throughputs
  - http://www.spec.org/jbb2000

Figure 9-51. SPEC JBB2000 Benchmark AM711.0

Notes:

SPECJBB2000 focuses to measure Java server performance for business applications. Java business benchmark (JBB) is helpful in predicting the performance and scalability of Java-based business solutions. It illustrates the Java engine effectiveness and how efficiently each processor, RAM, and disks are performing.

SPECJBB2000 benchmark highlights include:

- Emulation of a three-tier system, the most common type of server-size Java application today
- Business logic and object manipulation, the work of middle tier
- Clients replaced by driver threads
- Database storage by binary trees of objects
- Increasing amounts of workload applied, providing a graphical view of scalability

SPECweb99 Benchmark

- Measures Web serving ability
  - Static and dynamic page requests
  - Standardized workload
  - Measures simultaneous connections

- Metrics:
  - SPECweb99
  - http://www.spec.org/web99

Figure 9-52. SPECweb99 Benchmark AM711.0

Notes:

The SPECweb99 benchmark is designed to measure a system's ability to act as a Web server answering static and dynamic page requests. The features of SPECweb99 include:

- Standardized workload, agreed by major players in the world-wide Web (WWW) market
- Full disclosures available on this Web site
- Stable implementation with no incomparable versions
- Measurement of simultaneous connections rather than Hypertext Transfer Protocol (HTTP) operations

200 IBM @server pSeries Sizing and Capacity Planning

- Simulation of connections at a limited line speed
- Dynamic GETs, as well as static GETs; POST operations
- Keepalives (HTTP 1.0) and persistent connections (HTTP 1.1)
- Dynamic ad rotation using cookies and table lookups

- File accesses more closely matching today's real-world Web server access patterns
- An automated installation program for Microsoft Windows NT and UNIX installation scripts
- Inter-client communication using sockets

High Performance Computing Benchmarks

- SPEC CPU2000
  - Emphasizes performance of: CPU, memory, compilers
  - CINT2000 for integer performance
  - CFP2000 for floating point performance

- Metrics
  - SPECint2000/SPECint_rate2000
  - SPECint_base2000/SPECint_rate_base2000

Figure 9-53. High Performance Computing Benchmarks AM711.0

Notes:

SPEC CPU2000 is a benchmark that measures computer performance for CPU-intensive computing. SPEC CPU2000 contains two sets of benchmarks.

SPEC CPU2000 focuses on compute-intensive performance, which means these benchmarks emphasize the performance of:

- The computer processor (CPU)
- The memory architecture
- The compilers

SPEC CPU2000 is made up of two subcomponents that focus on two different types of compute intensive performance:

- CINT2000 for compute-intensive integer performance
- CFP2000 for compute-intensive floating point performance

SPEC CPU2000 is not intended to stress other system components, such as disk drives, networking, and graphics, which are not included in the benchmarks.

However, these components may affect a system configured in a particular way.

LINPACK Benchmark

- Linear Algebra Package library of subroutines
  - Three versions of benchmark
  - Used to determine the ‘Top 500 List’
- Metrics:
  - Millions of floating point operations per second (MFLOPS)
  - Billions of floating point operations per second (GFLOPS)
  - http://www.netlib.org/linpack

Figure 9-54. LINPACK Benchmark AM711.0

Notes:

LINear algebra PACKage (LINPACK) is the name of a library of subroutines for linear algebra calculations. It is also the name of a widely used benchmark that measures the performance of computers when solving a particular system of linear equations. The LINPACK library was superseded by the Linear Algebra PACKage (LAPACK) library. However, the name LINPACK still applies to the performance benchmark (even if the computations are done using the LAPACK library code). The LINPACK benchmark includes three versions of benchmarks:

- Linpack Fortran n=100 benchmark: This benchmark solves a 100 x 100 system of linear equations. The ground rules for running this benchmark are that you cannot make any changes to the Fortran code, not even to the comments. Only compiler optimization can be used to enhance performance.
- Linpack n=1000 benchmark1000 (also known as LINPACK Toward Peak Performance (TPP)). This benchmark solves a 1000 x 1000 system of linear equations. The ground rules for running this benchmark are more relaxed in that you can specify any linear

equation solve you want, implemented in any language. This is still a popular version of the benchmark for uniprocessors, but it is too small to be useful on parallel computers.

- LINPACK NxN (also known as LINPACK HPC or LINPACK Parallel. This benchmark solves an N x N system of linear equations. The value N can be quite large so that many processors can be applied with good parallel scalability. This benchmark attempts to measure the best performance of a system in solving a system of equations.

ISV Benchmarks

- SAP standard application benchmarks
  - mySAP business suite
  - Help test and verify scalability
  - Suite of benchmarks
  - http://www.sap.com/benchmark

Figure 9-55. ISV Benchmarks AM711.0

Notes:

SAP standard application benchmarks test and prove the scalability of the mySAP business suite. The benchmark results provide basic sizing recommendations for clients by testing new hardware, system software components, and RDBMS. They allow for the comparison of different system configurations.

The original SAP standard application benchmarks are available for many SAP components. Originally introduced to strengthen quality assurance, the SAP standard application benchmarks can help to test and verify scalability, concurrency and multi-user behavior of system software components, RDBMS, and business applications. All performance data relevant to system, user, and business applications are monitored during a benchmark run. They can help to compare platforms and offer basic input for sizing recommendations.

This section explains some popular benchmarks. For additional benchmark information, see: http://www.sap.com/benchmark.

Oracle Applications Standard Benchmark

- Focused on ERP applications
- Suite of benchmarks
  - Focus on mid-range business and a variety of Oracle applications
  - http://www.oracle.com/apps_benchmark/index.html

Figure 9-56. Oracle Applications Standard Benchmark AM711.0

Notes:

The Oracle applications standard benchmark (OASB) is a comparable standard workload. It demonstrates the performance and scalability of Oracle applications and provides metrics for the comparison of the performance of Oracle applications on different system configurations.

Chapter 4. Benchmarks **211**

The OASB is focused on ERP applications, simulating realistic client scenarios using a selection of the most commonly used Oracle applications modules.

Definitions of transactions that compose the benchmark load are obtained through collaboration with implementation consultants. They represent typical client workloads, with both OLTP and batch components.

The database used in the benchmark is designed to represent amounts of information that are typical of mid-market type businesses, whose annual revenues range from $100 million to $500 million US. This database is provided with the benchmark kit and is common to all platforms on which the benchmark is available.

Definitions of transactions that compose the OASB workload were obtained through collaboration with functional consultants. They represent typical client workloads, with batch transactions representing 25% of the total workload.

Siebel Platform Sizing and Performance Program

- Siebel 7 e-business suite
  - Aid in system planning and sizing
  - Simulates large organizations with several thousand users
  - Suite of programs
  - http://www.siebel.com/products/performance_benchmark/index.html

Figure 9-57. Siebel Platform Sizing and Performance Program AM711.0

Notes:

This benchmark provides performance and scalability characteristics of the Siebel 7 e-business Application Suite on various operating systems and database platforms. You can use this information as an aid for system planning and sizing to support business requirements.

The test simulates real-world requirements of a large organization with several thousand concurrent users from the call center (sales and service representatives), partner organizations (Partner Relationship Management), and clients (Web sales and Web service) community. It also simulates the requirements of users supporting application services such as work assignment (Siebel Assignment Manager) and business process management (Siebel Workflow). The test also simulates integration with legacy systems (Siebel EAI MQ Series Adapter) and Web systems (Siebel EAI HTTP Adapter).

Relative Performance Benchmark

- rPerf is an estimate of commercial processing performance
- Simulates CPU, cache, and memory operations
- Metric:
  - rPerf

Figure 9-58. Relative Performance Benchmark AM711.0

Notes:

Workloads have shifted over the last several years. IBM is committed to providing clients with a relative system performance metric that reflects those changes.

IBM publishes the rPerf measurement for the pSeries family of UNIX servers. It combines several different measures of total systems commercial performance and considers the demands on a Web server in today's environment.

rPerf is an estimate of commercial processing performance relative to other pSeries systems. It is derived from an IBM analytical model that uses characteristics from IBM internal workloads, TPC, and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be used in that way.

The model simulates such system operations as CPU, cache, and maximum memory available. The model does not simulate disk or network I/O operations.

Although the model uses general database and operating system parameters, the model does not reflect specific database or AIX version or release.

rPerf estimates are calculated based on systems with maximum memory, the latest levels of AIX 5L, and other pertinent software. Actual performance varies based on configuration details. The pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to compare estimated IBM UNIX commercial processing performance, actual system performance may vary.

It depends on many factors including system hardware configuration and software design and configuration.

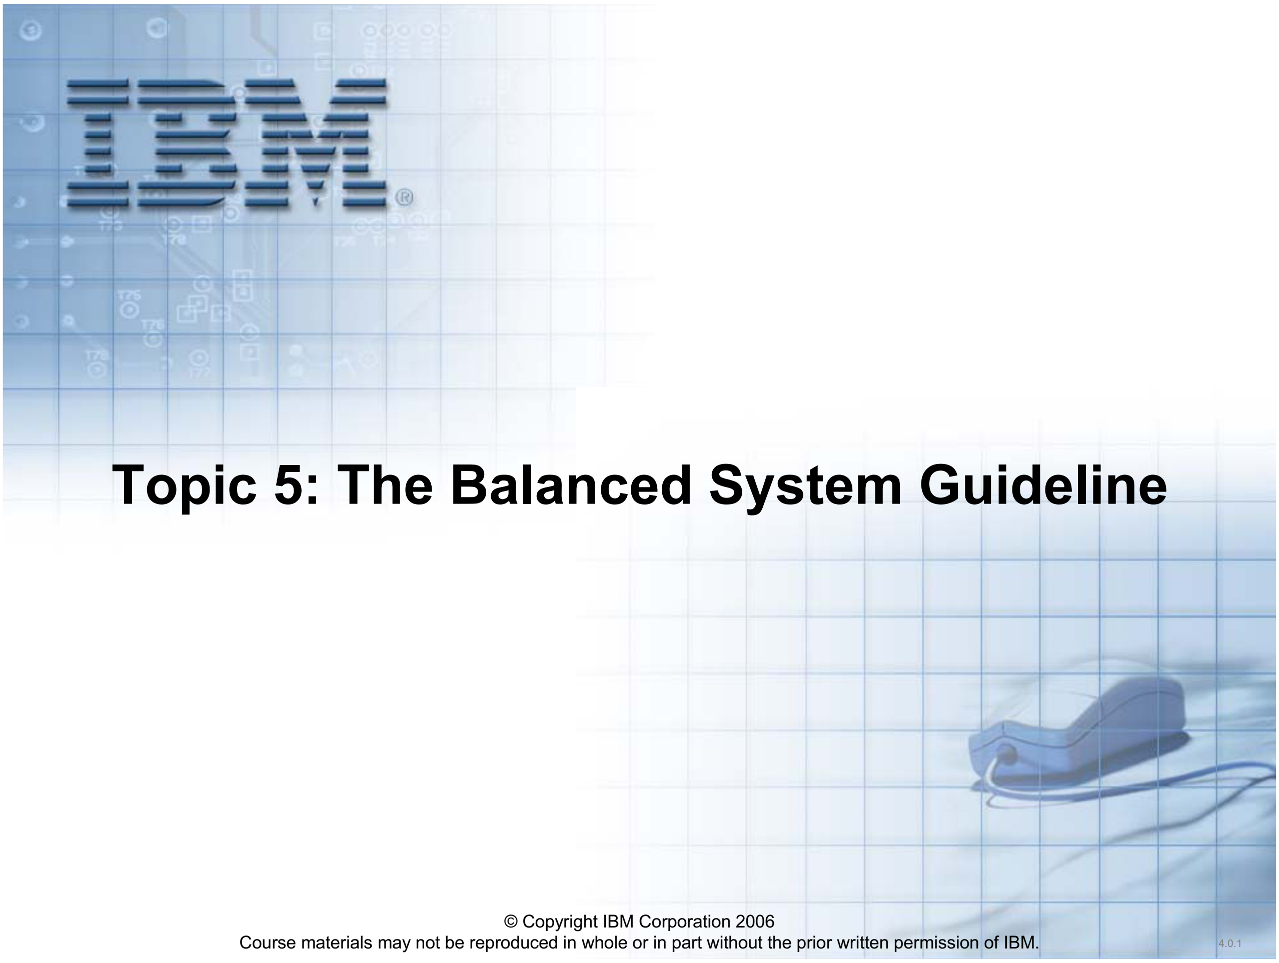

Topic 5: The Balanced System Guideline

4.0.1

Figure 9-59. Topic 5: The Balanced System Guideline AM711.0

Notes:

The Balanced System Guideline

- What is it?
  - Spreadsheet-based tool
- Why do I need it?
  - Lack of detail about requirements
  - Little information about workload
  - Variety of responses are required
- How can I get it?
  - IBM employees and Business Partners only

Figure 9-60. The Balanced System Guideline AM711.0

Notes:

In this unit we will begin our exploration of a very useful set of tools for system sizing. The "*Balanced System Guideline"* is a spreadsheet developed by IBM technical support over several years. A rule of thumb is an experience-based rule. It is better than a guess but not as precise as a result based on empirical data. To begin with we will discuss the underlying 'rules of thumb' that are used in the spreadsheet. If you are going to use the Balanced System Guideline tool it is better to be aware of its implicit assumptions.

When confronted with the need to generate system sizing without being provided with complete information by the client, the Balanced System Guideline lets you generate an estimate based on certain assumptions.

The Balanced System Guideline can be obtained from the following locations:

IBM employees go to the IBM System Sales internal Web site and search for BSG.

IBM Business Partners go to the Partnerinfo Web site and then to System Sales and then search for BSG.

In this course you will be provided with the spreadsheet on paper as a handout and a soft copy will be provided on the course CD. Please note that the copy provided with this course may not be the latest version and is for instructional purposes only. For real system sizing be sure to download the latest version of the tool.

Spreadsheet Pros and Cons

PRO

- ✓ Can be saved for future reference
- ✓ Can be sent to requesters
- ✓ Can be modified or used for ‘what if’ experiments
- ✓ Software to create it is widely available
- ✓ Looks professional and well thought out

CON

- ✗ Underlying spreadsheet data may be out of date
- ✗ Sense of accuracy may be greater than it actually is
- ✗ Inaccurate data or incorrect units lead to incorrect results
- ✗ Extensive use of rPerf CPU power rating may not fit the workload type

Figure 9-61. Spreadsheet Pros and Cons AM711.0

Notes:

The Balanced System Guideline spreadsheet should always be saved ‘as' to protect the original version of the spreadsheet tool. You can then have a collection of spreadsheets for different projects which provide a reference. The spreadsheet software is Excel 2002. The resulting spreadsheet can also be sent to requesters for their information for double checking or refinement. The nature of the spreadsheet means that the underlying assumptions, the ‘rules of thumb' are exposed to the end user. This may be seen as a pro or con. This is different though from a binary sizing tool where the assumptions are not accessible and only the result is visible. As new hardware models are introduced the spreadsheet will be updated. It is the responsibility of the user to ensure that they are using the latest version of the spreadsheet. It is also the case that information for a particular model will not be updated if performance increases may result from new releases of application software. One possible problem with the spreadsheet, according to the authors of the sizing Redbook is that too much data can be entered into the spreadsheet cells and data corruption can occur. Possibly the most subtle downside to the sizing tool is that it confers a heightened sense of accuracy on the result. Results which are created with three decimal place precision may actually be in the range of plus or minus 20% due to the nature of estimating. Other errors may be introduced into the result by incorrect data entry. Recognizing this type of error might be difficult if the person entering the data does not have a reasonable understanding of system sizing and can spot unlikely results. Remember, this is a tool not a magic wand. Finally, the CPU power rating which is a major calibration for the spreadsheet is derived in part from rPerf numbers which themselves are based on performance against a particular type of workload. If the workload for the system being sized does not fit the rPerf type then the results may not be applicable.

General Sizing

- Use 'rules of thumb' to estimate size of solution.

- Six golden principles of sizing:
  1. Correct processor configuration
  2. Balanced systems
  3. CPU magic number calculations
  4. Estimating CPU power
  5. Estimating memory size
  6. Estimating disk size

Figure 9-62. General Sizing AM711.0

Notes:

These principles have been incorporated into the Balanced System Guideline.

The processor configuration and sizing are critical. After you realize the configuration and sizing of the processors the other components are relatively simple to determine.

You must size all systems to yield a reasonably balanced system in regard to processor, memory, and disk. This allows for balance from the top down and the bottom up.

CPU magic number calculations, estimating CPU power and estimating memory sizing are all based on formulae which are part of the BSG. The formulae with examples are provided in this unit.

Estimating disk size can be done at a simple level in terms of capacity or by detailed measurement of the disk I/O rates for transactions either logical or physical. Sizing based on data size is becoming increasingly error prone due to the ever increasing disk sizes.

Correct Processor Configuration

- Determine the correct processor configuration first.
- Then configure other components to create a balanced system.

Figure 9-63. Correct Processor Configuration AM711.0

Notes:

Balanced Systems

- Processors must be powerful enough to make use of memory and keep disks busy.
- Disks must be fast enough to keep data supplied to memory.
- Memory must be able to feed main processor caches fast enough to keep the processors operating efficiently.

Figure 9-64. Balanced Systems AM711.0

Notes:

One of the underlying assumptions is that to perform well systems should be balanced with respect to CPU, memory and disks. In some cases 'unbalanced' systems may perform a particular workload better. If this is the case then the reason for the unbalanced design must be clearly stated in the design documentation.

CPU Magic Number Calculations

- How do I determine the correct ratio for a balanced system?
- What is the CPU magic number?

CPU seconds per transaction =

(CPU magic number)

$$\frac{\text{CPU rPerf x CPU Utilization x 3600}}{\text{Users x Transactions-per-hour-per-user}}$$

Figure 9-65. CPU Magic Number Calculations AM711.0

Notes:

The CPU magic number formula was derived from examination of production systems at IBM client sites and from testing in IBM benchmark centers around the world. An example is provided on the next slide. CPU seconds per transaction is defined as; The time it would take to finish a single transaction on a one rPerf machine. The one rPerf machine is defined as a 1 CPU p640 @ 375MHZ.

The magic number is a calibration factor which is used in the calculation of

CPU Magic Number Calculations

$$\text{CPU seconds per transaction} = \frac{\text{CPU rPerf x CPU Utilization x 3600}}{\text{Users x Transactions-per-hour-per-user}}$$

$$\text{CPU seconds per transaction} = \frac{10 \times 0.75 \times 3600}{1000 \times 60}$$

$$\text{CPU seconds per transaction} = 0.45$$

Figure 9-66. CPU Magic Number Calculations AM711.0

Notes:

In the example above 1000 users are performing 60 transactions per hour on a 10 rPerf system with a 75% CPU utilization. The result shows that it takes 4.5 seconds of processor time on the 10 rPerf system to run each transaction, therefore it would take a single rPerf system 0.45 seconds to run the same transaction.

The CPU utilization assumes the rule of thumb that CPU utilization is good when it is no more than 75% busy. This avoids using 100% of the power of the CPU.

Estimating CPU Power

- To calculate the estimated CPU power required for a workload use:

$$\frac{\text{Users x Transactions-per-hour-per-user x magic number}}{3600}$$

- Example: 1000 users at 60 transactions per hour

$$\text{Estimated CPU Power} = \frac{1000 \times 60 \times 0.45}{3600} = 7.5$$

Figure 9-67. Estimating CPU Power AM711.0

Notes:

In this calculation we use the calibration factor from the previous step, (the magic number) as part of the equation. The estimated CPU power is a measure of the power required for a particular workload. The example shows that a system with 1000 users performing 60 transactions per hour with a magic number of 0.45 would require a system with an estimated CPU power of 7.5. This shows how the 10 rPerf system from the previous step with a utilization of 75% would have an estimated power of 7.5. You should be able to see how the magic number builds in some margin. Choose a system from the pSeries models to provide the necessary 7.5 CPU power.

Estimating Memory Size

- Simple calculation gives memory size for different workloads and gives CPU power rating.

- More detailed level calculations are performed for each of the following memory uses:
  - Core system
  - Per user memory
  - Disks cache
  - Application binary size

Figure 9-68. Estimating Memory Size AM711.0

Notes:

Core system includes AIX, daemons, and basic file systems. It is estimated using a standard number.

Per user memory is the amount required to run user processes and batch tasks. Per user memory needs to be multiplied by the number of users on the system. Is best measured but can be guessed based on experience.

Disks cache is estimated for database like applications. It can be recommended by the vendor based on data volumes, measured on other systems of similar data volumes or estimated as a proportion of data volumes.

Application binary size should be measured.

Estimating Disk Size

- At a simple level the Balanced System Guideline indicates disk sizing for different workloads
- Disk sizing can be done first and the recommended CPU power rating can be generated
- Detailed level provides for measuring disk I/O rates for transactions

Figure 9-69. Estimating Disk Size AM711.0

Notes:

Balanced System Guideline - Introduction

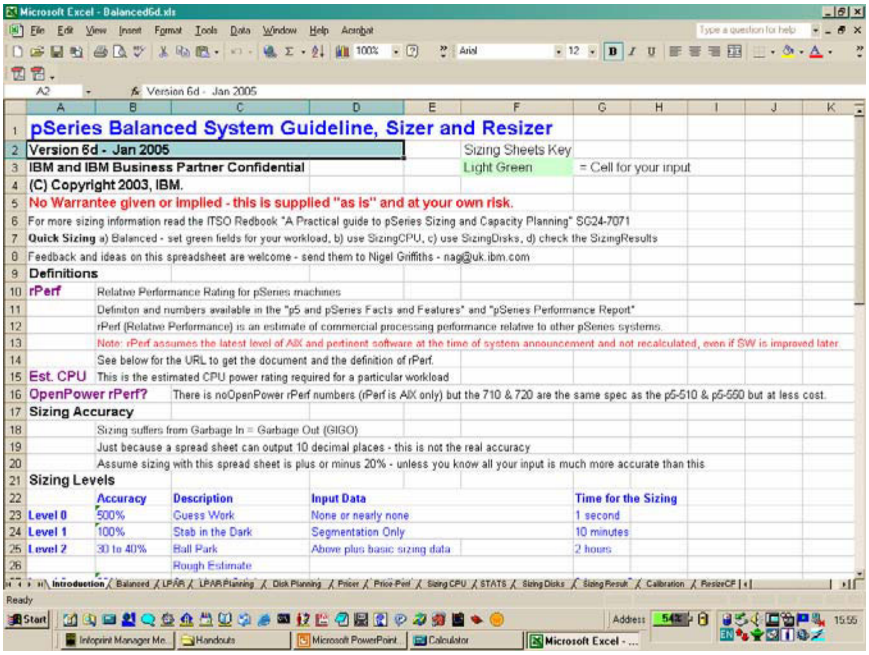

pSeries Balanced System Guideline, Sizer and Resizer

Version 6d - Jan 2005

Sizing Sheets Key

Light Green = Cell for your input

IBM and IBM Business Partner Confidential

(C) Copyright 2003, IBM.

No Warrantee given or implied - this is supplied "as is" and at your own risk.

For more sizing information read the ITSO Redbook "A Practical guide to pSeries Sizing and Capacity Planning" SG24-7071

Quick Sizing a) Balanced - set green fields for your workload, b) use SizingCPU, c) use SizingDisks, d) check the SizingResults

Feedback and ideas on this spreadsheet are welcome - send them to Nigel Griffiths - nag@uk.ibm.com

Definitions

rPerf Relative Performance Rating for pSeries machines

Definiton and numbers available in the "p5 and pSeries Facts and Features" and "pSeries Performance Report"

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other pSeries systems.

Note: rPerf assumes the latest level of AIX and pertinent software at the time of system announcement and not recalculated, even if SW is improved later

See below for the URL to get the document and the definition of rPerf.

Est. CPU This is the estimated CPU power rating required for a particular workload

OpenPower rPerf? There is noOpenPower rPerf numbers (rPerf is AIX only) but the 710 & 720 are the same spec as the p5-510 & p5-550 but at less cost.

Sizing Accuracy

Sizing suffers from Garbage In = Garbage Out (GIGO)

Just because a spread sheet can output 10 decimal places - this is not the real accuracy

Assume sizing with this spread sheet is plus or minus 20% - unless you know all your input is much more accurate than this

Sizing Levels

	Accuracy	**Description**	**Input Data**	**Time for the Sizing**
Level 0	500%	Guess Work	None or nearly none	1 second
Level 1	100%	Stab in the Dark	Segmentation Only	10 minutes
Level 2	30 to 40%	Ball Park	Above plus basic sizing data	2 hours
		Rough Estimate		

Figure 9-70. Balanced System Guideline - Introduction AM711.0

Notes:

The ‘Introduction' sheet contains version details, a few definitions and some useful hints and tips.

Balanced System Guideline - Balanced

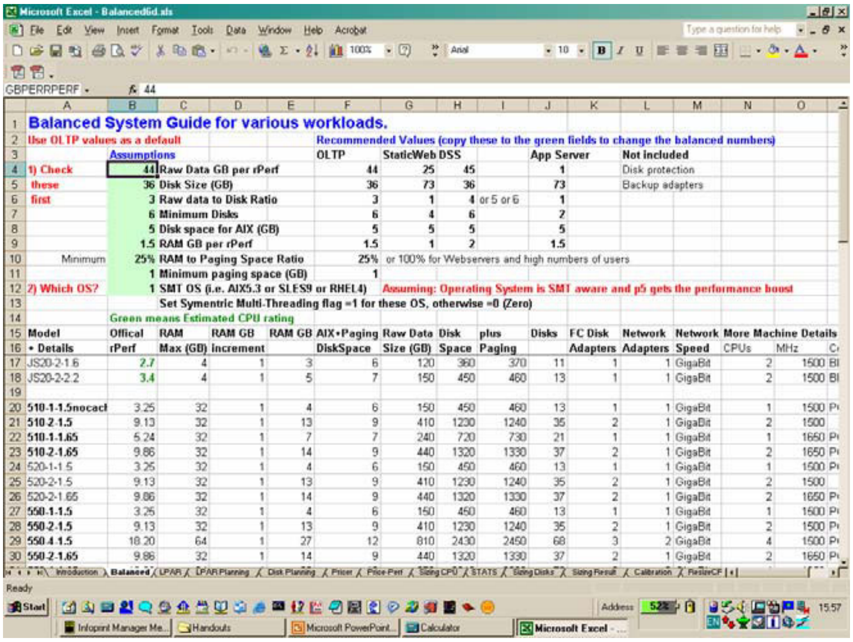

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	
1	Balanced System Guide for various workloads.															
2	Use OLTP values as a default					Recommended Values (copy these to the green fields to change the balanced numbers)										
3		Assumptions				OLTP	StaticWeb	DSS		App Server		Not included				
4	1) Check	44	Raw Data GB per rPerf			44	25	45		1		Disk protection				
5	these	36	Disk Size (GB)			36	73	36		73		Backup adapters				
6	first	3	Raw data to Disk Ratio			3	1	4	or 5 or 6	1						
7		6	Minimum Disks			6	4	6		2						
8		5	Disk space for AIX (GB)			5	5	5		5						
9		1.5	RAM GB per rPerf			1.5	1	2		1.5						
10	Minimum	25%	RAM to Paging Space Ratio			25%	or 100% for Webservers and high numbers of users									
11		1	Minimum paging space (GB)			1										
12	2) Which OS?	1	SMT OS (i.e. AIX5.3 or SLES9 or RHEL4)				Assuming: Operating System is SMT aware and p5 gets the performance boost									
13			Set Symentric Multi-Threading flag =1 for these OS, otherwise =0 (Zero)													
14		Green means Estimated CPU rating														
15	Model	Offical	RAM	RAM GB	RAM GB	AIX+Paging	Raw Data	Disk	plus	Disks	FC Disk	Network	Network	More Machine Details		
16	+ Details	rPerf	Max (GB)	Increment		DiskSpace	Size (GB)	Space	Paging		Adapters	Adapters	Speed	CPUs	MHz	C
17	JS20-2-1.6	2.7	4	1	3	6	120	360	370	11	1	1	GigaBit	2	1500	BI
18	JS20-2-2.2	3.4	4	1	5	7	150	450	460	13	1	1	GigaBit	2	1500	BI
19																
20	510-1-1.5nocacl	3.25	32	1	4	6	150	450	460	13	1	1	GigaBit	1	1500	P
21	510 2-1.5	9.13	32	1	13	9	410	1230	1240	35	2	1	GigaBit	2	1500	
22	510-1-1.65	5.24	32	1	7	7	240	720	730	21	1	1	GigaBit	1	1650	P
23	510-2-1.65	9.86	32	1	14	9	440	1320	1330	37	2	1	GigaBit	2	1650	P
24	520-1-1.5	3.25	32	1	4	6	150	450	460	13	1	1	GigaBit	1	1500	P
25	520-2-1.5	9.13	32	1	13	9	410	1230	1240	35	2	1	GigaBit	2	1500	
26	520-2-1.65	9.86	32	1	14	9	440	1320	1330	37	2	1	GigaBit	2	1650	P
27	550-1-1.5	3.25	32	1	4	6	150	450	460	13	1	1	GigaBit	1	1500	P
28	550-2-1.5	9.13	32	1	13	9	410	1230	1240	35	2	1	GigaBit	2	1500	P
29	550 4-1.5	18.20	64	1	27	12	810	2430	2450	68	3	2	GigaBit	4	1500	P
30	550 2-1.65	9.86	32	1	14	9	440	1320	1330	37	2	1	GigaBit	2	1650	P

Figure 9-71. Balanced System Guideline - Balanced AM711.0

Notes:

The ‘Balanced' sheet contains reference data about the current pSeries systems and the performance numbers. Also included are the ‘rules of thumb' used to suggest balanced systems. The workload characteristics can be modified here (the light green cells) so that it models Online Transaction Processing (OLTP), simple Web servers using Common Gateway Interface (CGI), Business Intelligence (BI) or application servers.

Balanced System Guideline - Balanced (1 of 2)

- Raw Data GB per rPerf
- Disk Size
- Raw Data to Disk Ratio
- Minimum Disks
- Disk Space for AIX
- RAM GB per rPerf
- RAM to Paging Space Ratio
- Minimum Paging Space

Figure 9-72. Balanced System Guideline - Balanced (1 of 2) AM711.0

Notes:

The top section of the Worksheet titled ‘Balanced' contains a series of assumptions with values associated with them in the light green filled cells. The assumptions are based on rules of thumb and can be modified if the workload requires it. There are also recommended values for other workloads; OLTP, Static Web Server, DSS and Application Server. The details of the assumptions are as follows.

Raw Data GB per rPerf: The number of GB of disk space that a system with a one rPerf rating is likely to keep busy. This is based on benchmark and production system analysis.

Disk Size: Should be set to the smallest or next smallest disk size available. Small numbers of large disks cannot support the same rates of disk I/O as large numbers of smaller ones. The field is changeable because disk technology moves rapidly and smaller disks become unavailable for purchase.

Raw Data to Disk Ratio: Normally set to three (1:3) for OLTP workloads which use an RDBMS.

Minimum Disks: For databases the recommendation is to set the minimum to six.

Disk Space for AIX: Recommendation is for 5GB but you can change it if required.

RAM GB per rPerf: Rule of thumb for amount of memory to allow per rPerf in the system for good performance. Older rules were for IGB or 2GB per CPU but as CPUs become faster these rules no longer apply as they are not based on the power of the processor.

RAM to Paging Space Ratio: With larger memory sizes, well controlled numbers of users, and AIX optimized paging space algorithms, this is a reasonable amount of paging space. If you can't control the number of users or vendors recommend more, you can use 100%. For AIX 4.3.3 add an additional 100%.

Minimum Paging Space: The value in GBs is the smallest amount of paging space allowed.

Balanced System Guideline - Balanced (2 of 2)

- Model
- rPerf
- RAM Max
- RAM GB Increment
- RAM GB Typical
- AIX + Paging Disk Space
- Raw Data Size
- Disk Space
- Plus Paging
- Disks
- Disk Adapters
- Network Adapters
- Network Speed
- Other fields

Figure 9-73. Balanced System Guideline - Balanced (2 of 2) AM711.0

Notes:

In the lower half of the worksheet the following column headings may be found. Each row contains the values for the model identified in the first column.

Model: The model column gives enough detail to make it clear which model and processor details are being referenced.

rPerf: This is the relative CPU performance rating used across the pSeries range.

RAM Max (GB): This is the maximum amount of memory that you can configure in the system in GBs.

RAM GB Typical: This is the recommended amount of memory. It is based on a simple rule of thumb in the GB RAM per rPerf cell at the top of the sheet and then rounds up the amount to a configured amount.

AIX + Paging Disk Space: This is the calculation of the disk space required to install AIX and paging space. When typical disks were around 2GB and memory was often below 1GB paging space was an important factor in sizing. Processors have become much faster, most systems have at least 1 GB of memory and disks are now generally greater in size than the paging requirements. In addition the AIX paging algorithms have improved and delayed allocation is now used. All these factors make the paging space requirements less important but they still cannot be ignored.

Raw Data Size: This is calculated from rPerf and the Raw Data GB per rPerf. This value represents the raw data volume for the database.

Disk Space: This value is calculated from the Raw Data size and the Data to Disk Ratio and factors in the extra space that a database needs.

Plus Paging: This value adds the AIX and paging space requirements to the previous column.

Disks: This value is calculated from the previous column and the preferred Disk Size. It is the number of recommended disks.

Disk Adapters: This value is calculated from the number of Disks column using a rule of thumb.

Network Adapters: This value is calculated from the rPerf column using a rule of thumb.

Network Speed: This value shows the network type. In most cases, this is now gigabit Ethernet which has become the standard for servers.

Other fields: More information about the models.

Balanced System Guideline - LPAR

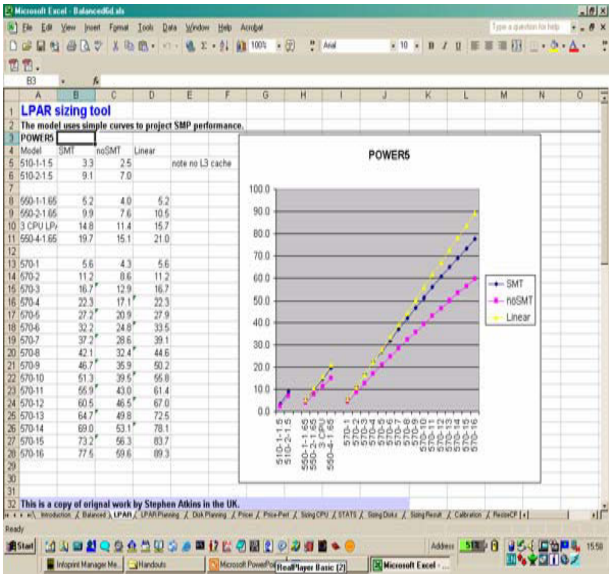

Microsoft Excel - BalancedGd.xls

	A	B	C	D	E
1	**LPAR sizing tool**				
2	**The model uses simple curves to project SMP performance.**				
3	**POWER5**				
4	Model	SMT	noSMT	Linear	
5	510-1-1.5	3.3	2.5		note no L3 cache
6	510-2-1.5	9.1	7.0		
7					
8	550-1-1.65	5.2	4.0	5.2	
9	550-2-1.65	9.9	7.6	10.5	
10	3 CPU LP/	14.8	11.4	15.7	
11	550-4-1.65	19.7	15.1	21.0	
12					
13	570-1	5.6	4.3	5.6	
14	570-2	11.2	8.6	11.2	
15	570-3	16.7	12.9	16.7	
16	570-4	22.3	17.1	22.3	
17	570-5	27.2	20.9	27.9	
18	570-6	32.2	24.8	33.5	
19	570-7	37.2	28.6	39.1	
20	570-8	42.1	32.4	44.6	
21	570-9	46.7	35.9	50.2	
22	570-10	51.3	39.5	55.8	
23	570-11	55.9	43.0	61.4	
24	570-12	60.5	46.5	67.0	
25	570-13	64.7	49.8	72.5	
26	570-14	69.0	53.1	78.1	
27	570-15	73.2	56.3	83.7	
28	570-16	77.5	59.6	89.3	
29					
30					
31					
32	**This is a copy of orignal work by Stephen Atkins in the UK.**				

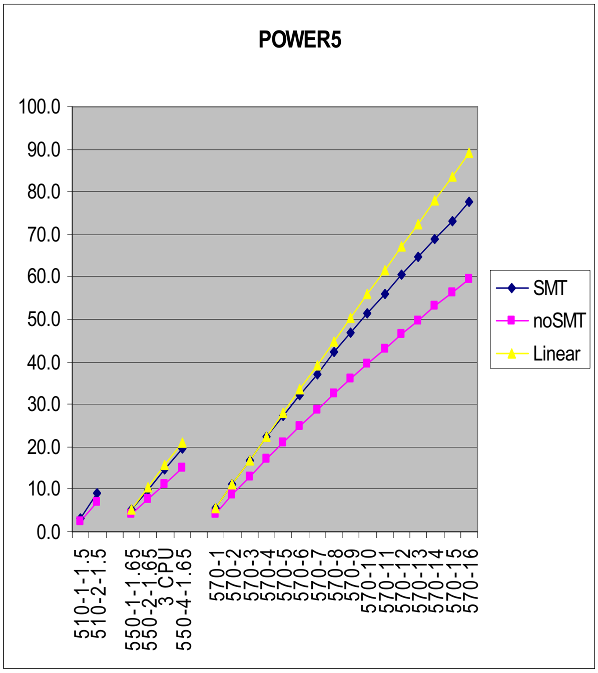

Figure 9-74. Balanced System Guideline - LPAR AM711.0

Notes:

The LPAR sheet provides all the information you need to perform sizing with LPARs so you can break up a larger system into smaller units. The sheet contains configuration information by system which provides an estimate of the rPerf value for a range of models and LPAR sizes. The results are graphed against linear scaling and show the result with and without SMT enabled.

Balanced System Guideline – LPAR Planning

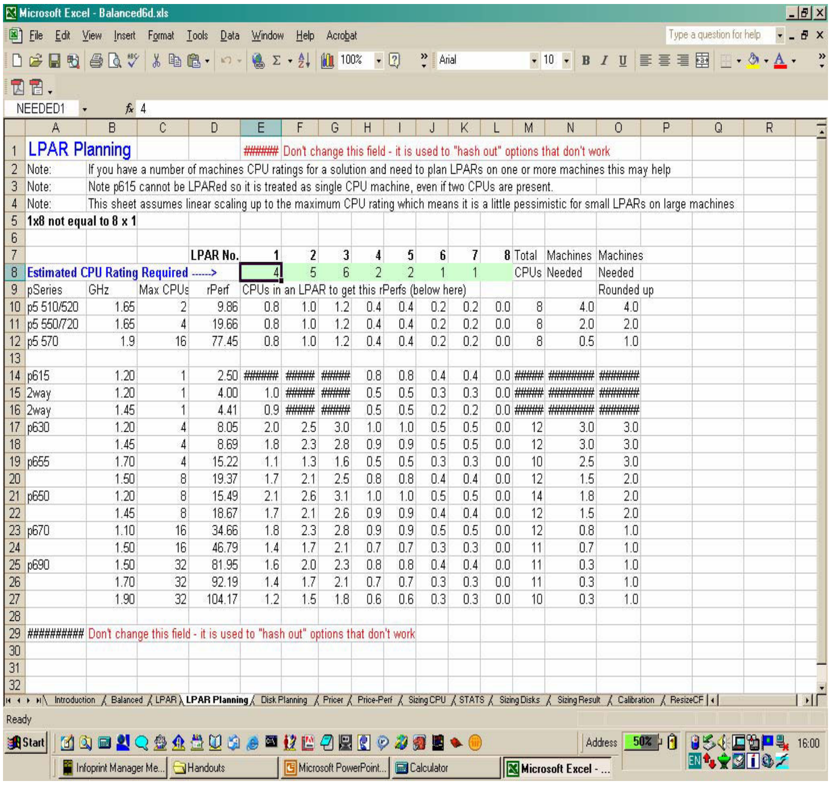

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1	LPAR Planning				######	Don't change this field - it is used to "hash out" options that don't work									
2	Note:	If you have a number of machines CPU ratings for a solution and need to plan LPARs on one or more machines this may help													
3	Note:	Note p615 cannot be LPARed so it is treated as single CPU machine, even if two CPUs are present.													
4	Note:	This sheet assumes linear scaling up to the maximum CPU rating which means it is a little pessimistic for small LPARs on large machines													
5	**1x8 not equal to 8 x 1**														
6															
7				**LPAR No.**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	Total	Machines	Machines
8	**Estimated CPU Rating Required ------>**				4	5	6	2	2	1	1		CPUs	Needed	Needed
9	pSeries	GHz	Max CPUs	rPerf	CPUs in an LPAR to get this rPerfs (below here)										Rounded up
10	p5 510/520	1.65	2	9.86	0.8	1.0	1.2	0.4	0.4	0.2	0.2	0.0	8	4.0	4.0
11	p5 550/720	1.65	4	19.66	0.8	1.0	1.2	0.4	0.4	0.2	0.2	0.0	8	2.0	2.0
12	p5 570	1.9	16	77.45	0.8	1.0	1.2	0.4	0.4	0.2	0.2	0.0	8	0.5	1.0
13															
14	p615	1.20	1	2.50	######	#####	#####	0.8	0.8	0.4	0.4	0.0	#####	########	#######
15	2way	1.20	1	4.00	1.0	#####	#####	0.5	0.5	0.3	0.3	0.0	#####	########	#######
16	2way	1.45	1	4.41	0.9	#####	#####	0.5	0.5	0.2	0.2	0.0	#####	########	#######
17	p630	1.20	4	8.05	2.0	2.5	3.0	1.0	1.0	0.5	0.5	0.0	12	3.0	3.0
18		1.45	4	8.69	1.8	2.3	2.8	0.9	0.9	0.5	0.5	0.0	12	3.0	3.0
19	p655	1.70	4	15.22	1.1	1.3	1.6	0.5	0.5	0.3	0.3	0.0	10	2.5	3.0
20		1.50	8	19.37	1.7	2.1	2.5	0.8	0.8	0.4	0.4	0.0	12	1.5	2.0
21	p650	1.20	8	15.49	2.1	2.6	3.1	1.0	1.0	0.5	0.5	0.0	14	1.8	2.0
22		1.45	8	18.67	1.7	2.1	2.6	0.9	0.9	0.4	0.4	0.0	12	1.5	2.0
23	p670	1.10	16	34.66	1.8	2.3	2.8	0.9	0.9	0.5	0.5	0.0	12	0.8	1.0
24		1.50	16	46.79	1.4	1.7	2.1	0.7	0.7	0.3	0.3	0.0	11	0.7	1.0
25	p690	1.50	32	81.95	1.6	2.0	2.3	0.8	0.8	0.4	0.4	0.0	11	0.3	1.0
26		1.70	32	92.19	1.4	1.7	2.1	0.7	0.7	0.3	0.3	0.0	11	0.3	1.0
27		1.90	32	104.17	1.2	1.5	1.8	0.6	0.6	0.3	0.3	0.0	10	0.3	1.0
28															
29	##########	Don't change this field - it is used to "hash out" options that don't work													
30															
31															
32															

Figure 9-75. Balanced System Guideline – LPAR Planning AM711.0

Notes:

Balanced System Guideline – Disk Planning

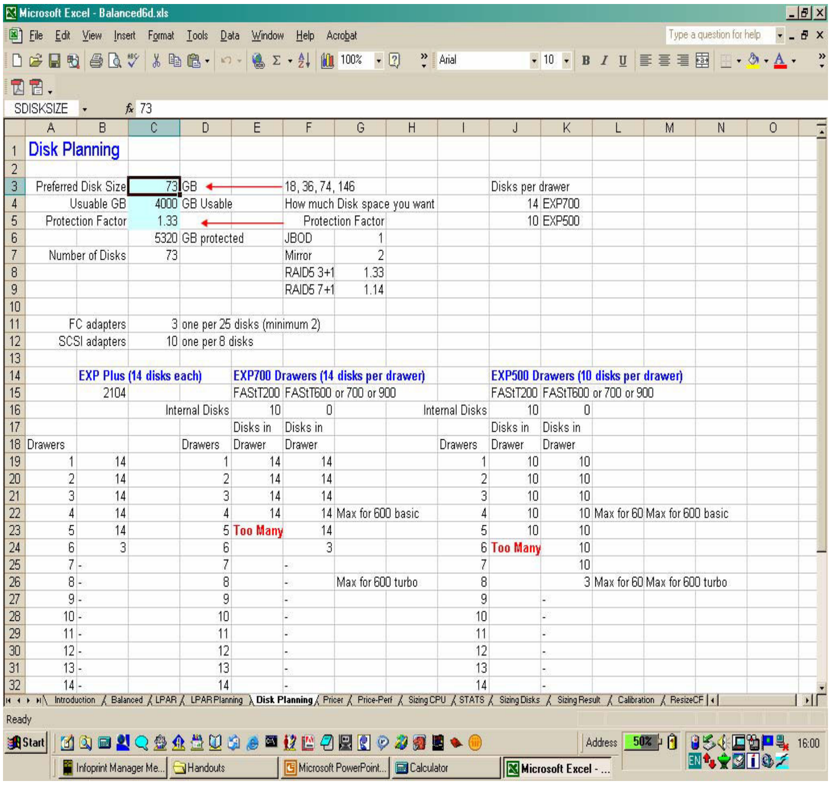

SDISKSIZE fx 73

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Disk Planning												
2													
3		Preferred Disk Size	73	GB		18, 36, 74, 146				Disks per drawer			
4		Usuable GB	4000	GB Usable		How much Disk space you want				14	EXP700		
5		Protection Factor	1.33			Protection Factor				10	EXP500		
6			5320	GB protected		JBOD	1						
7		Number of Disks	73			Mirror	2						
8						RAID5 3+1	1.33						
9						RAID5 7+1	1.14						
10													
11		FC adapters	3	one per 25 disks (minimum 2)									
12		SCSI adapters	10	one per 8 disks									
13													
14		EXP Plus (14 disks each)			EXP700 Drawers (14 disks per drawer)					EXP500 Drawers (10 disks per drawer)			
15		2104			FAStT200	FAStT600 or 700 or 900				FAStT200	FAStT600 or 700 or 900		
16				Internal Disks	10	0			Internal Disks	10	0		
17					Disks in	Disks in				Disks in	Disks in		
18	Drawers			Drawers	Drawer	Drawer			Drawers	Drawer	Drawer		
19	1	14		1	14	14			1	10	10		
20	2	14		2	14	14			2	10	10		
21	3	14		3	14	14			3	10	10		
22	4	14		4	14	14	Max for 600 basic		4	10	10	Max for 60	Max for 600 basic
23	5	14		5	Too Many	14			5	10	10		
24	6	3		6		3			6	Too Many	10		
25	7	-		7		-			7		10		
26	8	-		8		-	Max for 600 turbo		8		3	Max for 60	Max for 600 turbo
27	9	-		9		-			9		-		
28	10	-		10		-			10		-		
29	11	-		11		-			11		-		
30	12	-		12		-			12		-		
31	13	-		13		-			13		-		
32	14	-		14		-			14		-		

Figure 9-76. Balanced System Guideline – Disk Planning AM711.0

Notes:

Balanced System Guideline - Pricing

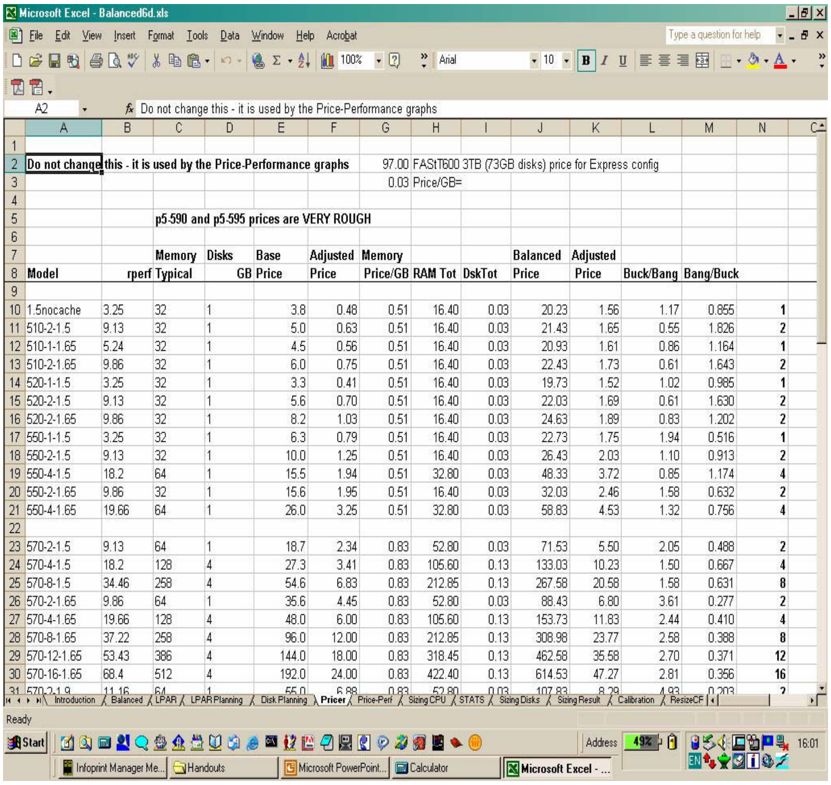

Do not change this - it is used by the Price-Performance graphs

97.00 FAStT600 3TB (73GB disks) price for Express config

0.03 Price/GB=

p5-590 and p5-595 prices are VERY ROUGH

Model	rperf	Memory Typical	Disks GB	Base Price	Adjusted Price	Memory Price/GB	RAM Tot	DskTot	Balanced Price	Adjusted Price	Buck/Bang	Bang/Buck	
1.5nocache	3.25	32	1	3.8	0.48	0.51	16.40	0.03	20.23	1.56	1.17	0.855	1
510-2-1.5	9.13	32	1	5.0	0.63	0.51	16.40	0.03	21.43	1.65	0.55	1.826	2
510-1-1.65	5.24	32	1	4.5	0.56	0.51	16.40	0.03	20.93	1.61	0.86	1.164	1
510-2-1.65	9.86	32	1	6.0	0.75	0.51	16.40	0.03	22.43	1.73	0.61	1.643	2
520-1-1.5	3.25	32	1	3.3	0.41	0.51	16.40	0.03	19.73	1.52	1.02	0.985	1
520-2-1.5	9.13	32	1	5.6	0.70	0.51	16.40	0.03	22.03	1.69	0.61	1.630	2
520-2-1.65	9.86	32	1	8.2	1.03	0.51	16.40	0.03	24.63	1.89	0.83	1.202	2
550-1-1.5	3.25	32	1	6.3	0.79	0.51	16.40	0.03	22.73	1.75	1.94	0.516	1
550-2-1.5	9.13	32	1	10.0	1.25	0.51	16.40	0.03	26.43	2.03	1.10	0.913	2
550-4-1.5	18.2	64	1	15.5	1.94	0.51	32.80	0.03	48.33	3.72	0.85	1.174	4
550-2-1.65	9.86	32	1	15.6	1.95	0.51	16.40	0.03	32.03	2.46	1.58	0.632	2
550-4-1.65	19.66	64	1	26.0	3.25	0.51	32.80	0.03	58.83	4.53	1.32	0.756	4
570-2-1.5	9.13	64	1	18.7	2.34	0.83	52.80	0.03	71.53	5.50	2.05	0.488	2
570-4-1.5	18.2	128	4	27.3	3.41	0.83	105.60	0.13	133.03	10.23	1.50	0.667	4
570-8-1.5	34.46	258	4	54.6	6.83	0.83	212.85	0.13	267.58	20.58	1.58	0.631	8
570-2-1.65	9.86	64	1	35.6	4.45	0.83	52.80	0.03	88.43	6.80	3.61	0.277	2
570-4-1.65	19.66	128	4	48.0	6.00	0.83	105.60	0.13	153.73	11.83	2.44	0.410	4
570-8-1.65	37.22	258	4	96.0	12.00	0.83	212.85	0.13	308.98	23.77	2.58	0.388	8
570-12-1.65	53.43	386	4	144.0	18.00	0.83	318.45	0.13	462.58	35.58	2.70	0.371	12
570-16-1.65	68.4	512	4	192.0	24.00	0.83	422.40	0.13	614.53	47.27	2.81	0.356	16
570-2-1.9	11.16	64	1	55.0	6.88	0.83	52.80	0.03	107.83	8.29	4.93	0.203	2

Figure 9-77. Balanced System Guideline - Pricing AM711.0

Notes:

Balanced System Guideline - Price/Performance

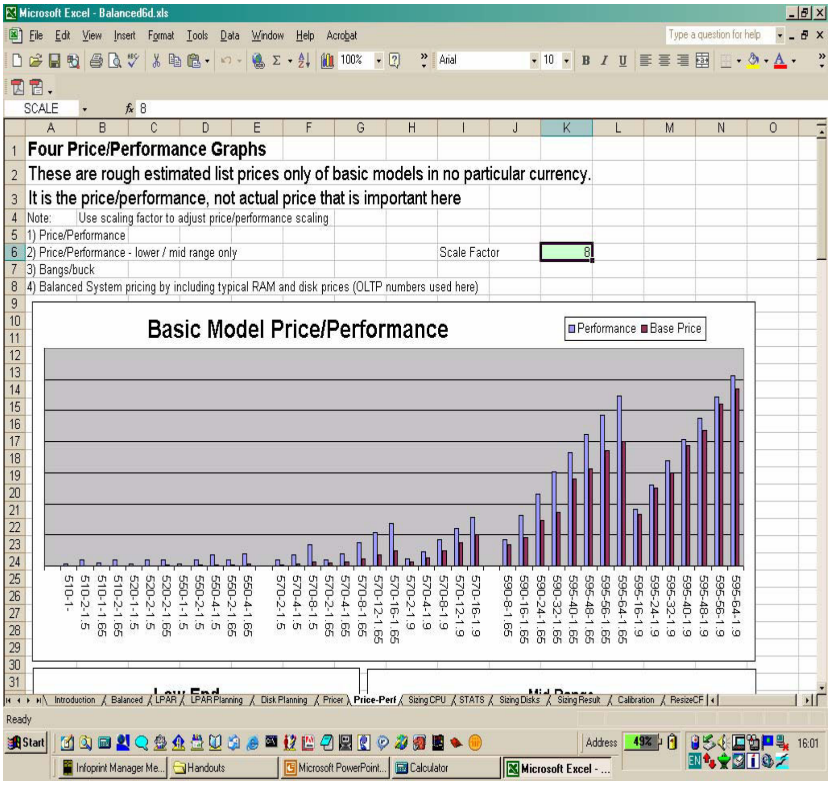

Four Price/Performance Graphs

These are rough estimated list prices only of basic models in no particular currency.

It is the price/performance, not actual price that is important here

Note: Use scaling factor to adjust price/performance scaling

1) Price/Performance

2) Price/Performance - lower / mid range only

Scale Factor: 8

3) Bangs/buck

4) Balanced System pricing by including typical RAM and disk prices (OLTP numbers used here)

Figure 9-78. Balanced System Guideline - Price/Performance AM711.0

Notes:

Balanced System Guideline – Sizing CPU

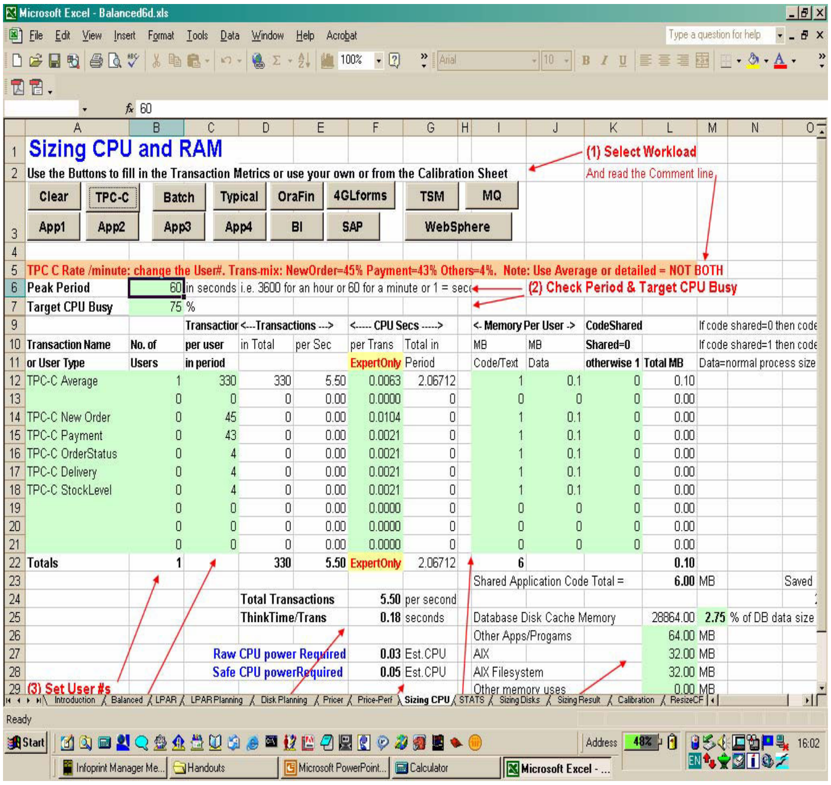

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1	Sizing CPU and RAM										(1) Select Workload				
2	Use the Buttons to fill in the Transaction Metrics or use your own or from the Calibration Sheet										And read the Comment line				
3															
4															
5	TPC C Rate /minute: change the User#. Trans-mix: NewOrder=45% Payment=43% Others=4%. Note: Use Average or detailed = NOT BOTH														
6	Peak Period	60	in seconds	i.e. 3600 for an hour or 60 for a minute or 1 = sec						(2) Check Period & Target CPU Busy					
7	Target CPU Busy	75	%												
9			Transactior	<---Transactions --->		<----- CPU Secs ----->			<- Memory Per User ->		CodeShared		If code shared=0 then code		
10	Transaction Name	No. of	per user	in Total	per Sec	per Trans	Total in		MB	MB	Shared=0		If code shared=1 then code		
11	or User Type	Users	in period			ExpertOnly	Period		Code/Text	Data	otherwise 1	Total MB	Data=normal process size		
12	TPC-C Average	1	330	330	5.50	0.0063	2.06712		1	0.1	0	0.10			
13		0	0	0	0.00	0.0000	0		0	0	0	0.00			
14	TPC-C New Order	0	45	0	0.00	0.0104	0		1	0.1	0	0.00			
15	TPC-C Payment	0	43	0	0.00	0.0021	0		1	0.1	0	0.00			
16	TPC-C OrderStatus	0	4	0	0.00	0.0021	0		1	0.1	0	0.00			
17	TPC-C Delivery	0	4	0	0.00	0.0021	0		1	0.1	0	0.00			
18	TPC-C StockLevel	0	4	0	0.00	0.0021	0		1	0.1	0	0.00			
19		0	0	0	0.00	0.0000	0		0	0	0	0.00			
20		0	0	0	0.00	0.0000	0		0	0	0	0.00			
21		0	0	0	0.00	0.0000	0		0	0	0	0.00			
22	Totals	1		330	5.50	ExpertOnly	2.06712		6			0.10			
23									Shared Application Code Total =			6.00	MB		Saved
24				Total Transactions		5.50	per second								
25				ThinkTime/Trans		0.18	seconds		Database Disk Cache Memory			28864.00	2.75	% of DB data size	
26									Other Apps/Progams			64.00	MB		
27			Raw CPU power Required			0.03	Est.CPU		AIX			32.00	MB		
28			Safe CPU powerRequired			0.05	Est.CPU		AIX Filesystem			32.00	MB		
29	(3) Set User #s								Other memory uses			0.00	MB		

© Copyright IBM Corporation 2006

Figure 9-79. Balanced System Guideline – Sizing CPU AM711.0

Notes:

Balanced System Guideline – Stats

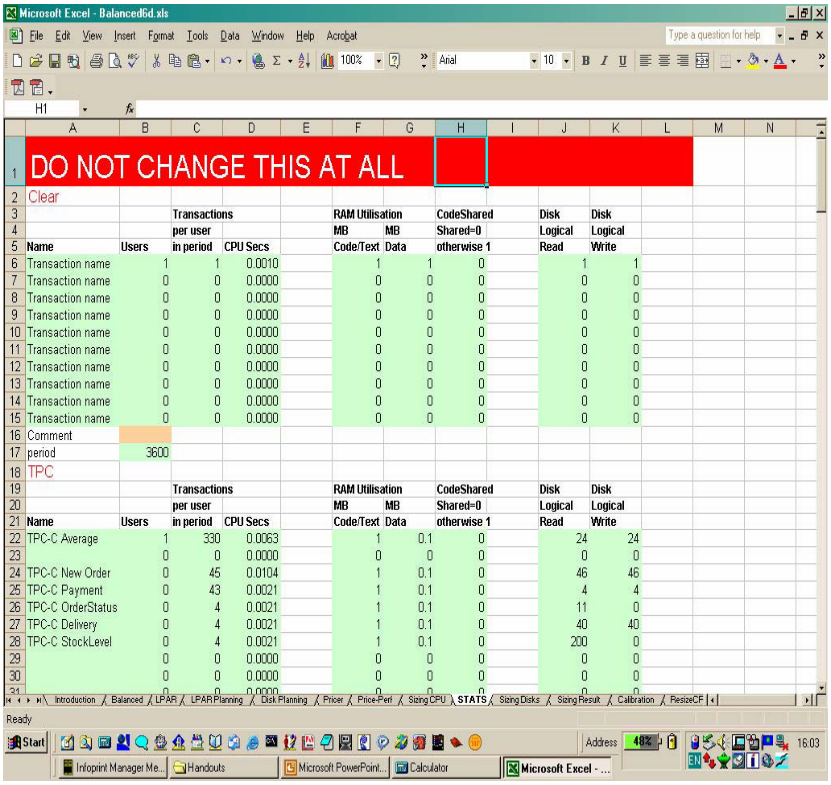

	A	B	C	D	E	F	G	H	I	J	K
1	DO NOT CHANGE THIS AT ALL										
2	Clear										
3			Transactions			RAM Utilisation		CodeShared		Disk	Disk
4			per user			MB	MB	Shared=0		Logical	Logical
5	Name	Users	in period	CPU Secs		Code/Text	Data	otherwise 1		Read	Write
6	Transaction name	1	1	0.0010		1	1	0		1	1
7	Transaction name	0	0	0.0000		0	0	0		0	0
8	Transaction name	0	0	0.0000		0	0	0		0	0
9	Transaction name	0	0	0.0000		0	0	0		0	0
10	Transaction name	0	0	0.0000		0	0	0		0	0
11	Transaction name	0	0	0.0000		0	0	0		0	0
12	Transaction name	0	0	0.0000		0	0	0		0	0
13	Transaction name	0	0	0.0000		0	0	0		0	0
14	Transaction name	0	0	0.0000		0	0	0		0	0
15	Transaction name	0	0	0.0000		0	0	0		0	0
16	Comment										
17	period	3600									
18	TPC										
19			Transactions			RAM Utilisation		CodeShared		Disk	Disk
20			per user			MB	MB	Shared=0		Logical	Logical
21	Name	Users	in period	CPU Secs		Code/Text	Data	otherwise 1		Read	Write
22	TPC-C Average	1	330	0.0063		1	0.1	0		24	24
23		0	0	0.0000		0	0	0		0	0
24	TPC-C New Order	0	45	0.0104		1	0.1	0		46	46
25	TPC-C Payment	0	43	0.0021		1	0.1	0		4	4
26	TPC-C OrderStatus	0	4	0.0021		1	0.1	0		11	0
27	TPC-C Delivery	0	4	0.0021		1	0.1	0		40	40
28	TPC-C StockLevel	0	4	0.0021		1	0.1	0		200	0
29		0	0	0.0000		0	0	0		0	0
30		0	0	0.0000		0	0	0		0	0

Figure 9-80. Balanced System Guideline – Stats AM711.0

Notes:

Balanced System Guideline – Sizing Disks

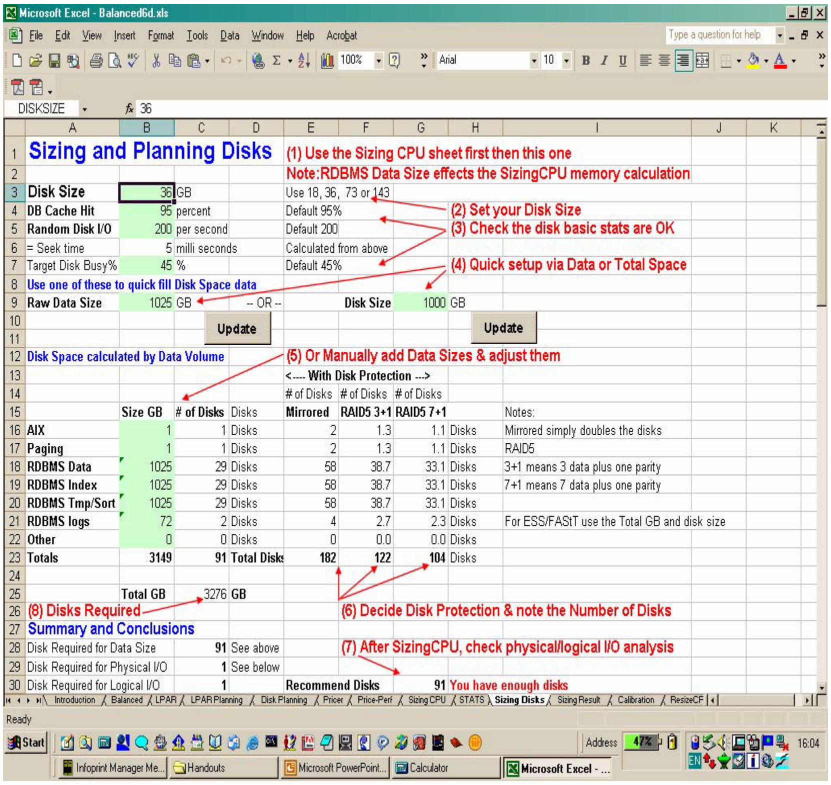

DISKSIZE fx 36

	A	B	C	D	E	F	G	H	I
1	Sizing and Planning Disks				(1) Use the Sizing CPU sheet first then this one				
2					Note:RDBMS Data Size effects the SizingCPU memory calculation				
3	**Disk Size**	36	GB		Use 18, 36, 73 or 143				
4	**DB Cache Hit**	95	percent		Default 95%		(2) Set your Disk Size		
5	**Random Disk I/O**	200	per second		Default 200		(3) Check the disk basic stats are OK		
6	= Seek time	5	milli seconds		Calculated from above				
7	Target Disk Busy%	45	%		Default 45%		(4) Quick setup via Data or Total Space		
8	Use one of these to quick fill Disk Space data								
9	**Raw Data Size**	1025	GB	-- OR --		**Disk Size**	1000	GB	
10			Update					Update	
11									
12	Disk Space calculated by Data Volume				(5) Or Manually add Data Sizes & adjust them				
13					<---- **With Disk Protection** --->				
14					# of Disks	# of Disks	# of Disks		
15		**Size GB**	**# of Disks**	Disks	**Mirrored**	**RAID5 3+1**	**RAID5 7+1**		Notes:
16	**AIX**	1	1	Disks	2	1.3	1.1	Disks	Mirrored simply doubles the disks
17	**Paging**	1	1	Disks	2	1.3	1.1	Disks	RAID5
18	**RDBMS Data**	1025	29	Disks	58	38.7	33.1	Disks	3+1 means 3 data plus one parity
19	**RDBMS Index**	1025	29	Disks	58	38.7	33.1	Disks	7+1 means 7 data plus one parity
20	**RDBMS Tmp/Sort**	1025	29	Disks	58	38.7	33.1	Disks	
21	**RDBMS logs**	72	2	Disks	4	2.7	2.3	Disks	For ESS/FAStT use the Total GB and disk size
22	**Other**	0	0	Disks	0	0.0	0.0	Disks	
23	**Totals**	**3149**	**91**	**Total Disks**	**182**	**122**	**104**	Disks	
24									
25		**Total GB**	3276	**GB**					
26	(8) Disks Required					(6) Decide Disk Protection & note the Number of Disks			
27	Summary and Conclusions								
28	Disk Required for Data Size		**91**	See above		(7) After SizingCPU, check physical/logical I/O analysis			
29	Disk Required for Physical I/O		**1**	See below					
30	Disk Required for Logical I/O		**1**		**Recommend Disks**		**91**	You have enough disks	

Figure 9-81. Balanced System Guideline – Sizing Disks

AM711.0

Notes:

Balanced System Guideline – Sizing Results

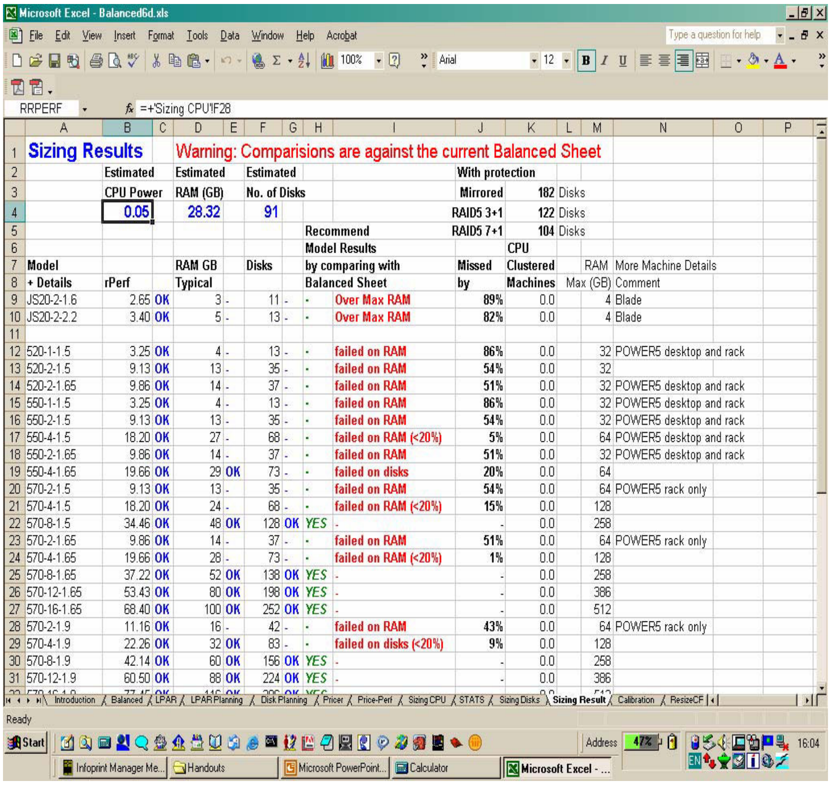

RRPERF =+'Sizing CPU'!F28

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
1	Sizing Results			Warning: Comparisions are against the current Balanced Sheet										
2		Estimated		Estimated		Estimated				With protection				
3		CPU Power		RAM (GB)		No. of Disks				Mirrored	182	Disks		
4		0.05		28.32		91				RAID5 3+1	122	Disks		
5								Recommend		RAID5 7+1	104	Disks		
6								Model Results			CPU			
7	Model			RAM GB		Disks		by comparing with		Missed	Clustered		RAM	More Machine Details
8	+ Details	rPerf		Typical				Balanced Sheet		by	Machines		Max (GB)	Comment
9	JS20-2-1.6	2.65	OK	3	-	11	-	·	Over Max RAM	89%	0.0		4	Blade
10	JS20-2-2.2	3.40	OK	5	-	13	-	·	Over Max RAM	82%	0.0		4	Blade
11														
12	520-1-1.5	3.25	OK	4	-	13	-	·	failed on RAM	86%	0.0		32	POWER5 desktop and rack
13	520-2-1.5	9.13	OK	13	-	35	-	·	failed on RAM	54%	0.0		32	
14	520-2-1.65	9.86	OK	14	-	37	-	·	failed on RAM	51%	0.0		32	POWER5 desktop and rack
15	550-1-1.5	3.25	OK	4	-	13	-	·	failed on RAM	86%	0.0		32	POWER5 desktop and rack
16	550-2-1.5	9.13	OK	13	-	35	-	·	failed on RAM	54%	0.0		32	POWER5 desktop and rack
17	550-4-1.5	18.20	OK	27	-	68	-	·	failed on RAM (<20%)	5%	0.0		64	POWER5 desktop and rack
18	550-2-1.65	9.86	OK	14	-	37	-	·	failed on RAM	51%	0.0		32	POWER5 desktop and rack
19	550-4-1.65	19.66	OK	29	OK	73	-	·	failed on disks	20%	0.0		64	
20	570-2-1.5	9.13	OK	13	-	35	-	·	failed on RAM	54%	0.0		64	POWER5 rack only
21	570-4-1.5	18.20	OK	24	-	68	-	·	failed on RAM (<20%)	15%	0.0		128	
22	570-8-1.5	34.46	OK	48	OK	128	OK	YES	-	-	0.0		258	
23	570-2-1.65	9.86	OK	14	-	37	-	·	failed on RAM	51%	0.0		64	POWER5 rack only
24	570-4-1.65	19.66	OK	28	-	73	-	·	failed on RAM (<20%)	1%	0.0		128	
25	570-8-1.65	37.22	OK	52	OK	138	OK	YES	-	-	0.0		258	
26	570-12-1.65	53.43	OK	80	OK	198	OK	YES	-	-	0.0		386	
27	570-16-1.65	68.40	OK	100	OK	252	OK	YES	-	-	0.0		512	
28	570-2-1.9	11.16	OK	16	-	42	-	·	failed on RAM	43%	0.0		64	POWER5 rack only
29	570-4-1.9	22.26	OK	32	OK	83	-	·	failed on disks (<20%)	9%	0.0		128	
30	570-8-1.9	42.14	OK	60	OK	156	OK	YES	-	-	0.0		258	
31	570-12-1.9	60.50	OK	88	OK	224	OK	YES	-	-	0.0		386	

Figure 9-82. Balanced System Guideline – Sizing Results AM711.0

Notes:

Balanced System Guideline – Calibration

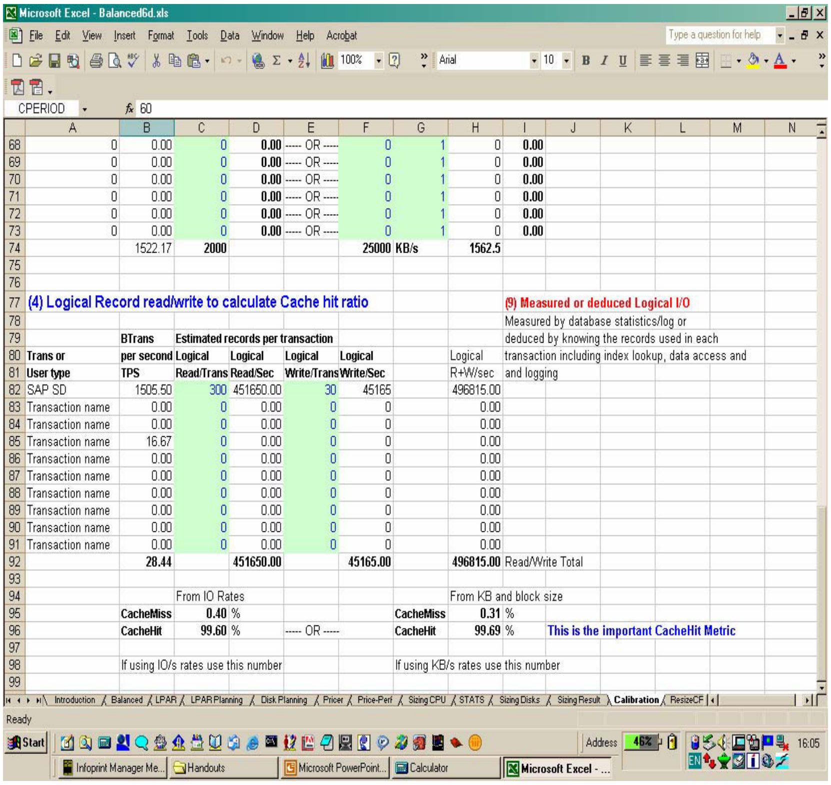

CPERIOD fx 60

	A	B	C	D	E	F	G	H	I	J	K	L	M	N
68	0	0.00	0	**0.00**	----- OR -----	0	1	0	**0.00**					
69	0	0.00	0	**0.00**	----- OR -----	0	1	0	**0.00**					
70	0	0.00	0	**0.00**	----- OR -----	0	1	0	**0.00**					
71	0	0.00	0	**0.00**	----- OR -----	0	1	0	**0.00**					
72	0	0.00	0	**0.00**	----- OR -----	0	1	0	**0.00**					
73	0	0.00	0	**0.00**	----- OR -----	0	1	0	**0.00**					
74		1522.17	**2000**			**25000**	**KB/s**	**1562.5**						
75														
76														
77	(4) Logical Record read/write to calculate Cache hit ratio								**(9) Measured or deduced Logical I/O**					
78									Measured by database statistics/log or					
79		**BTrans**	**Estimated records per transaction**						deduced by knowing the records used in each					
80	**Trans or**	**per second**	**Logical**	**Logical**	**Logical**	**Logical**		Logical	transaction including index lookup, data access and					
81	**User type**	**TPS**	**Read/Trans**	**Read/Sec**	**Write/Trans**	**Write/Sec**		R+W/sec	and logging					
82	SAP SD	1505.50	300	451650.00	30	45165		496815.00						
83	Transaction name	0.00	0	0.00	0	0		0.00						
84	Transaction name	0.00	0	0.00	0	0		0.00						
85	Transaction name	16.67	0	0.00	0	0		0.00						
86	Transaction name	0.00	0	0.00	0	0		0.00						
87	Transaction name	0.00	0	0.00	0	0		0.00						
88	Transaction name	0.00	0	0.00	0	0		0.00						
89	Transaction name	0.00	0	0.00	0	0		0.00						
90	Transaction name	0.00	0	0.00	0	0		0.00						
91	Transaction name	0.00	0	0.00	0	0		0.00						
92		**28.44**		**451650.00**		**45165.00**		**496815.00**	Read/Write Total					
93														
94			From IO Rates					From KB and block size						
95		**CacheMiss**	**0.40**	%			**CacheMiss**	**0.31**	%					
96		**CacheHit**	**99.60**	%	----- OR -----		**CacheHit**	**99.69**	%	**This is the important CacheHit Metric**				
97														
98		If using IO/s rates use this number					If using KB/s rates use this number							
99														

© Copyright IBM Corporation 2006

Figure 9-83. Balanced System Guideline – Calibration AM711.0

Notes:

Balanced System Guideline – Resize CPU

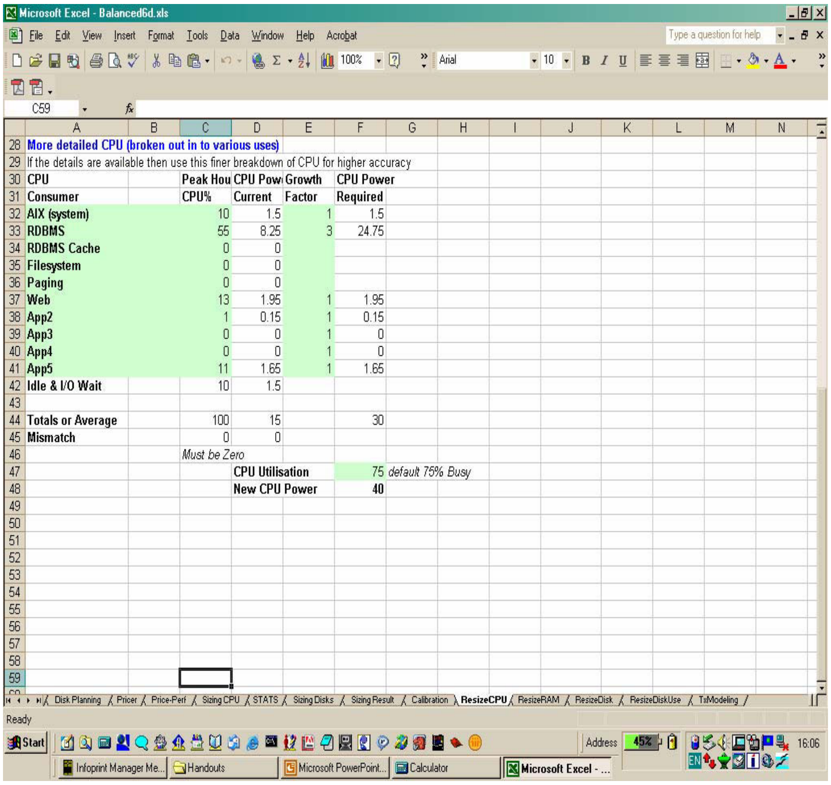

	A	B	C	D	E	F	G
28	More detailed CPU (broken out in to various uses)						
29	If the details are available then use this finer breakdown of CPU for higher accuracy						
30	CPU		Peak Hou	CPU Pow	Growth	CPU Power	
31	Consumer		CPU%	Current	Factor	Required	
32	AIX (system)		10	1.5	1	1.5	
33	RDBMS		55	8.25	3	24.75	
34	RDBMS Cache		0	0			
35	Filesystem		0	0			
36	Paging		0	0			
37	Web		13	1.95	1	1.95	
38	App2		1	0.15	1	0.15	
39	App3		0	0	1	0	
40	App4		0	0	1	0	
41	App5		11	1.65	1	1.65	
42	Idle & I/O Wait		10	1.5			
43							
44	Totals or Average		100	15		30	
45	Mismatch		0	0			
46			*Must be Zero*				
47				CPU Utilisation		75	*default 75% Busy*
48				New CPU Power		**40**	

© Copyright IBM Corporation 2006

Figure 9-84. Balanced System Guideline – Resize CPU AM711.0

Notes:

Balanced System Guideline – Resize RAM

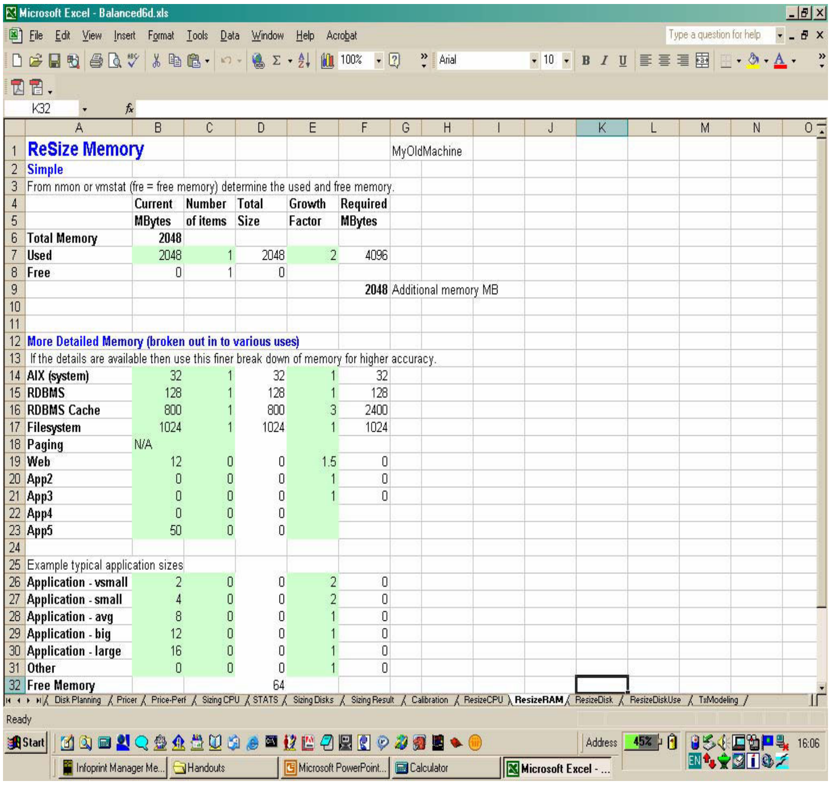

	A	B	C	D	E	F	G	H
1	**ReSize Memory**						MyOldMachine	
2	**Simple**							
3	From nmon or vmstat (fre = free memory) determine the used and free memory.							
4		**Current**	**Number**	**Total**	**Growth**	**Required**		
5		**MBytes**	**of items**	**Size**	**Factor**	**MBytes**		
6	**Total Memory**	**2048**						
7	**Used**	2048	1	2048	2	4096		
8	**Free**	0	1	0				
9						**2048**	Additional memory MB	
10								
11								
12	**More Detailed Memory (broken out in to various uses)**							
13	If the details are available then use this finer break down of memory for higher accuracy.							
14	**AIX (system)**	32	1	32	1	32		
15	**RDBMS**	128	1	128	1	128		
16	**RDBMS Cache**	800	1	800	3	2400		
17	**Filesystem**	1024	1	1024	1	1024		
18	**Paging**	N/A						
19	**Web**	12	0	0	1.5	0		
20	**App2**	0	0	0	1	0		
21	**App3**	0	0	0	1	0		
22	**App4**	0	0	0				
23	**App5**	50	0	0				
24								
25	Example typical application sizes							
26	**Application - vsmall**	2	0	0	2	0		
27	**Application - small**	4	0	0	2	0		
28	**Application - avg**	8	0	0	1	0		
29	**Application - big**	12	0	0	1	0		
30	**Application - large**	16	0	0	1	0		
31	**Other**	0	0	0	1	0		
32	**Free Memory**			64				

Figure 9-85. Balanced System Guideline – Resize RAM AM711.0

Notes:

Balanced System Guideline – Resize Disks

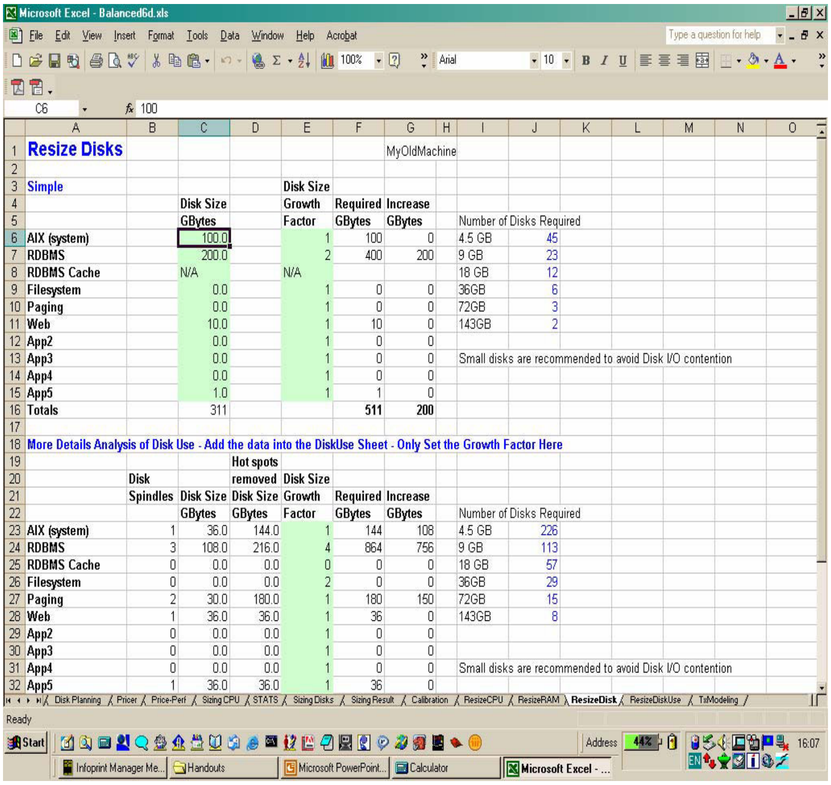

	A	B	C	D	E	F	G	H	I	J
1	**Resize Disks**						MyOldMachine			
2										
3	**Simple**				**Disk Size**					
4			**Disk Size**		**Growth**	**Required**	**Increase**			
5			**GBytes**		**Factor**	**GBytes**	**GBytes**		Number of Disks Required	
6	**AIX (system)**		100.0		1	100	0		4.5 GB	45
7	**RDBMS**		200.0		2	400	200		9 GB	23
8	**RDBMS Cache**		N/A		N/A				18 GB	12
9	**Filesystem**		0.0		1	0	0		36GB	6
10	**Paging**		0.0		1	0	0		72GB	3
11	**Web**		10.0		1	10	0		143GB	2
12	**App2**		0.0		1	0	0			
13	**App3**		0.0		1	0	0		Small disks are recommended to avoid Disk I/O contention	
14	**App4**		0.0		1	0	0			
15	**App5**		1.0		1	1	0			
16	**Totals**		311			**511**	**200**			
17										
18	**More Details Analysis of Disk Use - Add the data into the DiskUse Sheet - Only Set the Growth Factor Here**									
19				**Hot spots**						
20		**Disk**		**removed**	**Disk Size**					
21		**Spindles**	**Disk Size**	**Disk Size**	**Growth**	**Required**	**Increase**			
22			**GBytes**	**GBytes**	**Factor**	**GBytes**	**GBytes**		Number of Disks Required	
23	**AIX (system)**	1	36.0	144.0	1	144	108		4.5 GB	226
24	**RDBMS**	3	108.0	216.0	4	864	756		9 GB	113
25	**RDBMS Cache**	0	0.0	0.0	0	0	0		18 GB	57
26	**Filesystem**	0	0.0	0.0	2	0	0		36GB	29
27	**Paging**	2	30.0	180.0	1	180	150		72GB	15
28	**Web**	1	36.0	36.0	1	36	0		143GB	8
29	**App2**	0	0.0	0.0	1	0	0			
30	**App3**	0	0.0	0.0	1	0	0			
31	**App4**	0	0.0	0.0	1	0	0		Small disks are recommended to avoid Disk I/O contention	
32	**App5**	1	36.0	36.0	1	36	0			

Figure 9-86. Balanced System Guideline – Resize Disks AM711.0

Notes:

Balanced System Guideline – Resize Disk Use

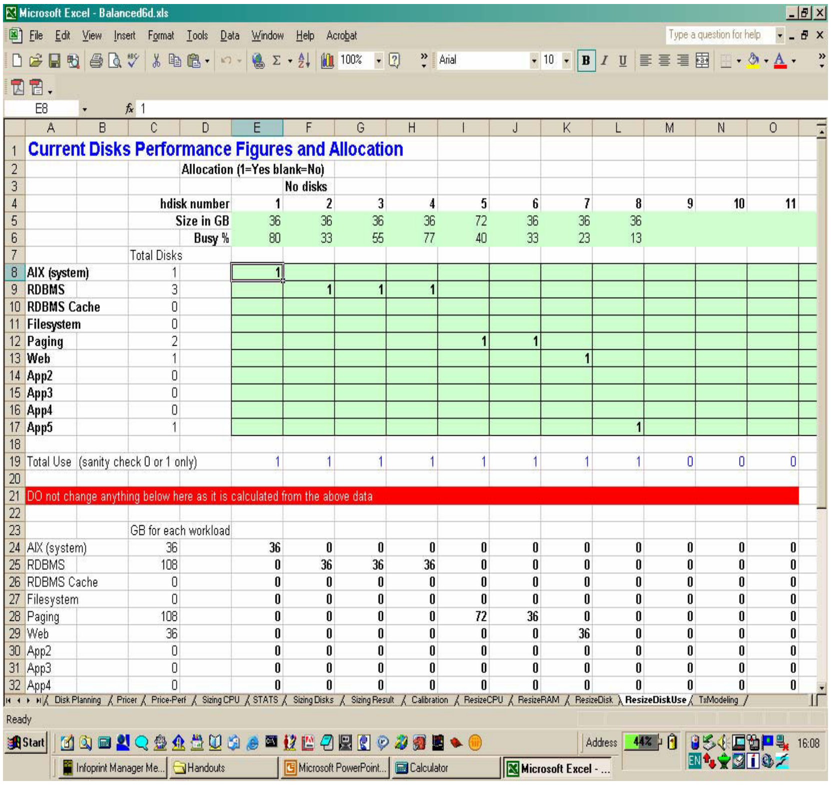

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
1	**Current Disks Performance Figures and Allocation**														
2				**Allocation (1=Yes blank=No)**											
3						**No disks**									
4				**hdisk number**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**
5				**Size in GB**	36	36	36	36	72	36	36	36			
6				**Busy %**	80	33	55	77	40	33	23	13			
7			Total Disks												
8	**AIX (system)**		1		**1**										
9	**RDBMS**		3			**1**	**1**	**1**							
10	**RDBMS Cache**		0												
11	**Filesystem**		0												
12	**Paging**		2						**1**	**1**					
13	**Web**		1								**1**				
14	**App2**		0												
15	**App3**		0												
16	**App4**		0												
17	**App5**		1									**1**			
18															
19	Total Use	(sanity check 0 or 1 only)			1	1	1	1	1	1	1	1	0	0	0
20															
21	DO not change anything below here as it is calculated from the above data														
22															
23			GB for each workload												
24	AIX (system)		36		**36**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
25	RDBMS		108		**0**	**36**	**36**	**36**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
26	RDBMS Cache		0		**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
27	Filesystem		0		**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
28	Paging		108		**0**	**0**	**0**	**0**	**72**	**36**	**0**	**0**	**0**	**0**	**0**
29	Web		36		**0**	**0**	**0**	**0**	**0**	**0**	**36**	**0**	**0**	**0**	**0**
30	App2		0		**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
31	App3		0		**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
32	App4		0		**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**

Figure 9-87. Balanced System Guideline – Resize Disk Use AM711.0

Notes:

Balanced System Guideline – Tx Modelling

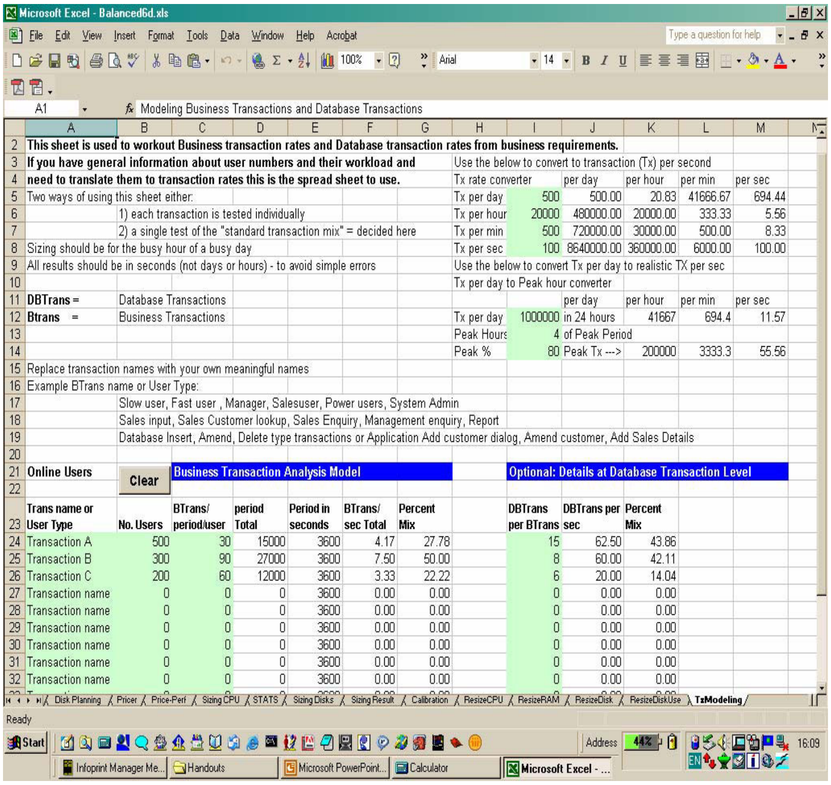

Modeling Business Transactions and Database Transactions

This sheet is used to workout Business transaction rates and Database transaction rates from business requirements.

If you have general information about user numbers and their workload and need to translate them to transaction rates this is the spread sheet to use.

Two ways of using this sheet either:

1) each transaction is tested individually

2) a single test of the "standard transaction mix" = decided here

Sizing should be for the busy hour of a busy day

All results should be in seconds (not days or hours) - to avoid simple errors

DBTrans = Database Transactions

Btrans = Business Transactions

Use the below to convert to transaction (Tx) per second

Tx rate converter		per day	per hour	per min	per sec
Tx per day	500	500.00	20.83	41666.67	694.44
Tx per hour	20000	480000.00	20000.00	333.33	5.56
Tx per min	500	720000.00	30000.00	500.00	8.33
Tx per sec	100	8640000.00	360000.00	6000.00	100.00

Use the below to convert Tx per day to realistic TX per sec

Tx per day to Peak hour converter

		per day	per hour	per min	per sec
Tx per day	1000000	in 24 hours	41667	694.4	11.57
Peak Hours	4	of Peak Period			
Peak %	80	Peak Tx --->	200000	3333.3	55.56

Replace transaction names with your own meaningful names

Example BTrans name or User Type:

Slow user, Fast user , Manager, Salesuser, Power users, System Admin

Sales input, Sales Customer lookup, Sales Enquiry, Management enquiry, Report

Database Insert, Amend, Delete type transactions or Application Add customer dialog, Amend customer, Add Sales Details

Online Users

Business Transaction Analysis Model

Optional: Details at Database Transaction Level

Trans name or User Type	**No. Users**	**BTrans/ period/user**	**period Total**	**Period in seconds**	**BTrans/ sec Total**	**Percent Mix**	**DBTrans per BTrans**	**DBTrans per sec**	**Percent Mix**
Transaction A	500	30	15000	3600	4.17	27.78	15	62.50	43.86
Transaction B	300	90	27000	3600	7.50	50.00	8	60.00	42.11
Transaction C	200	60	12000	3600	3.33	22.22	6	20.00	14.04
Transaction name	0	0	0	3600	0.00	0.00	0	0.00	0.00
Transaction name	0	0	0	3600	0.00	0.00	0	0.00	0.00
Transaction name	0	0	0	3600	0.00	0.00	0	0.00	0.00
Transaction name	0	0	0	3600	0.00	0.00	0	0.00	0.00
Transaction name	0	0	0	3600	0.00	0.00	0	0.00	0.00
Transaction name	0	0	0	3600	0.00	0.00	0	0.00	0.00

Figure 9-88. Balanced System Guideline – Tx Modelling AM711.0

Notes:

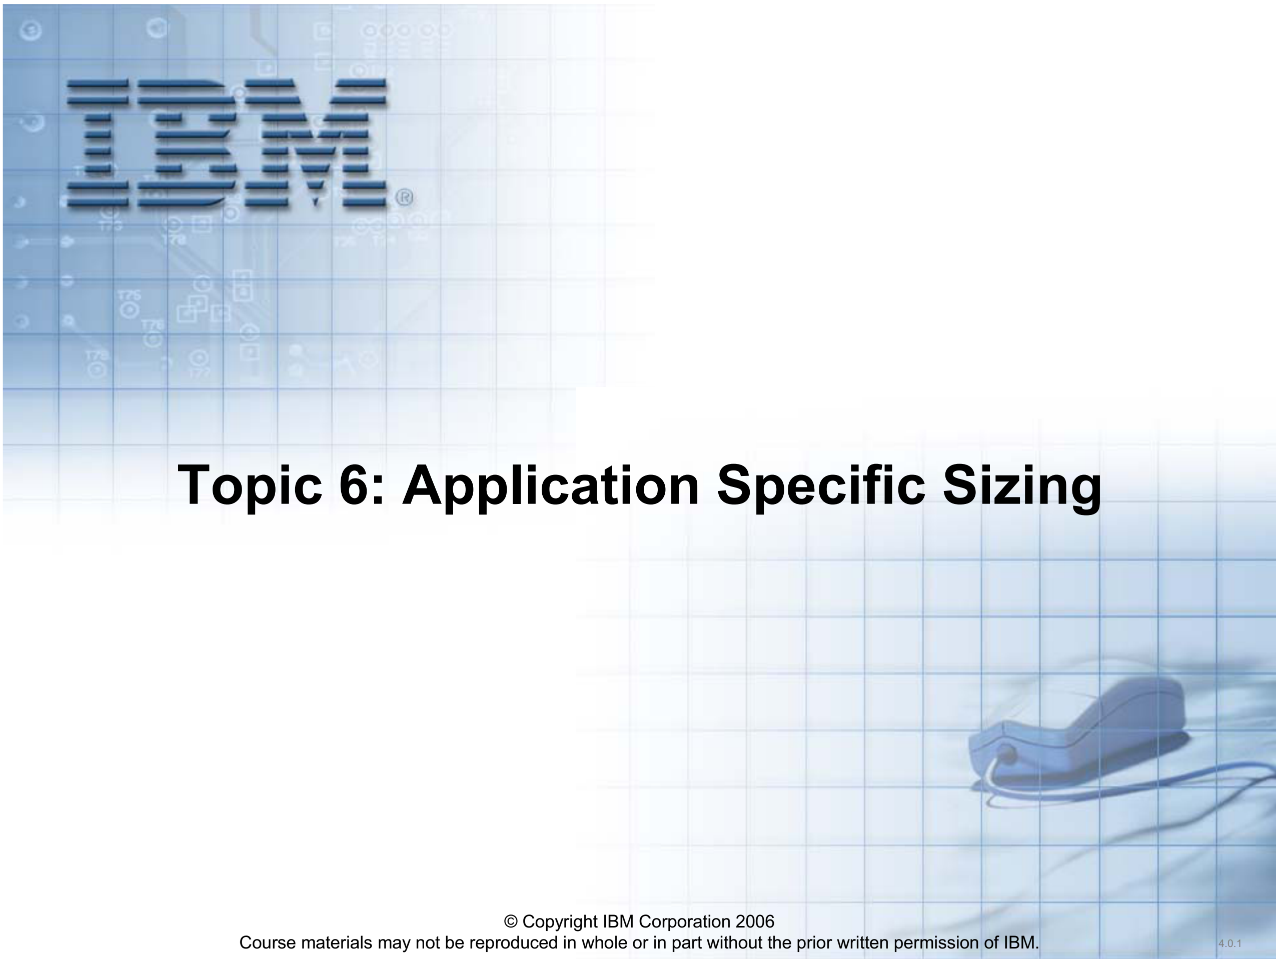

Topic 6: Application Specific Sizing

4.0.1

Figure 9-89. Topic 6: Application Specific Sizing AM711.0

Notes:

Application-Specific Sizing

- IBM Applications
  - DB2
  - Lotus Domino
  - Tivoli Storage Manager
  - WebSphere
- ISV Applications
  - eSizing@us.ibm.com sizing support
  - Quick e-sizing guides
  - IBM eServer Sizing Guide

Figure 9-90. Application-Specific Sizing AM711.0

Notes:

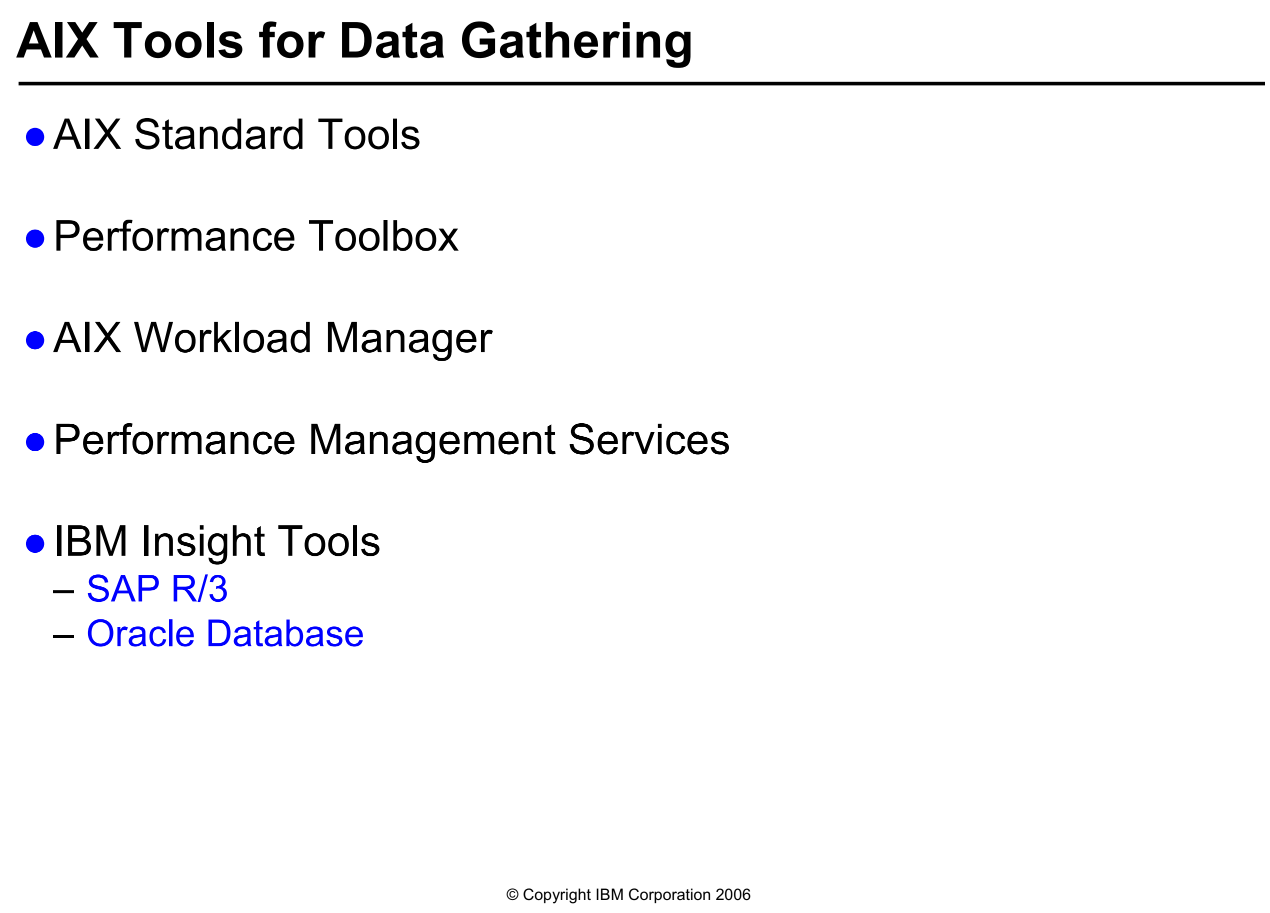

AIX Tools for Data Gathering

- AIX Standard Tools
- Performance Toolbox
- AIX Workload Manager
- Performance Management Services
- IBM Insight Tools
  - SAP R/3
  - Oracle Database

Figure 9-91. AIX Tools for Data Gathering AM711.0

Notes:

e-Config and Solution Design

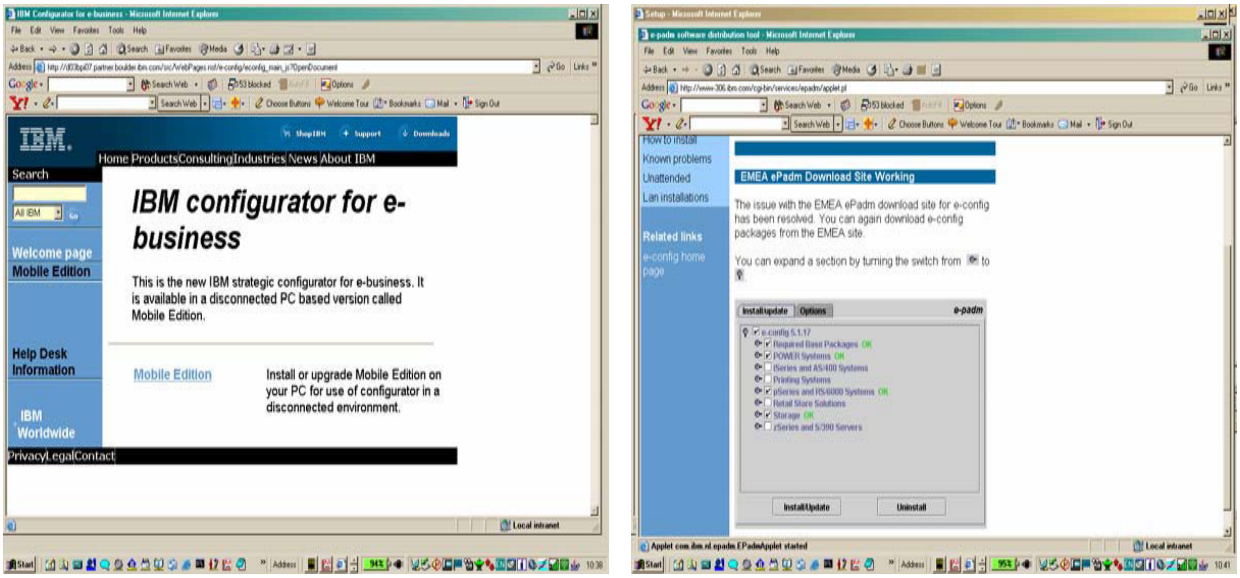
IBM Configurator for e-business - Microsoft Internet Explorer
http://d03bp07.partner.boulder.ibm.com/src/WebPages.nsf/e-config/econfig_main_js?OpenDocument
Shop IBM
Support
Downloads
Home Products Consulting Industries News About IBM
Search
All IBM
Welcome page
Mobile Edition
Help Desk Information
IBM Worldwide
Privacy Legal Contact
Local intranet
e-padm software distribution tool - Microsoft Internet Explorer
http://www-306.ibm.com/cgi-bin/services/epadm/applet.pl
How to install
Known problems
Unattended
Lan installations
Related links
e-config home page
Install/update
Options
e-padm
Install/Update
Uninstall
Applet com.ibm.nl.epadm.EPadmApplet started
Local intranet

IBM configurator for e-business

This is the new IBM strategic configurator for e-business. It is available in a disconnected PC based version called Mobile Edition.

Mobile Edition — Install or upgrade Mobile Edition on your PC for use of configurator in a disconnected environment.

EMEA ePadm Download Site Working

The issue with the EMEA ePadm download site for e-config has been resolved. You can again download e-config packages from the EMEA site.

You can expand a section by turning the switch from ⊶ to ⚲.

- ☑ e-config 5.1.17
  - ☑ Required Base Packages OK
  - ☑ POWER Systems OK
  - ☐ iSeries and AS/400 Systems
  - ☐ Printing Systems
  - ☑ pSeries and RS/6000 Systems OK
  - ☐ Retail Store Solutions
  - ☑ Storage OK
  - ☐ zSeries and S/390 Servers

Figure 9-92. e-Config and Solution Design AM711.0

Notes:

This class is not intended to teach you how to use the e-config tool but if you already know about it then you can use it as a sizing tool for the solution.

System Planning Tool (SPT)

- What is it?
- Where do I get it?
- How do I install it?
- How do I use it?

Figure 9-93. System Planning Tool (SPT)

Notes:

SPT Installation

- Download code from Web site
  - http://www-03.ibm.com/servers/eserver/support/tools/systemplanningtool/
- Installs on Windows 2000/XP
- Run the install wizard
- Read the user guide

Figure 9-94. SPT Installation AM711.0

Notes:

Checkpoint Questions (1 of 2)

1. Usually the main obstacle to accurate sizing is:
   a. Inaccurate tools
   b. Too much information
   c. Not enough information
   d. Lack of interest in accurate sizing

2. The rules of thumb for sizing are based on:
   a. Measurements taken on the client’s system
   b. Experience
   c. Pure guesswork
   d. Astrology

Figure 9-95. Checkpoint Questions (1 of 2) AM711.0

Notes:

Checkpoint Questions (2 of 2)

3. Which of the following questions is useful in segmenting a sizing problem?
   a. What is the type of application?
   b. What is the user doing? (workload type)
   c. How much data is being processed?
   d. All of the above

4. If the run queue length is greater than 10 times the number of processors in the system then, from a utilization point of view this is:
   a. Good
   b. Bad
   c. Ugly
   d. Normal

Figure 9-96. Checkpoint Questions (2 of 2) AM711.0

Notes:

Any Questions?

Figure 9-97. Any Questions? AM711.0

Notes:

Unit Summary

You should now:

- Understand the constraints placed on accurate sizing.
- Be able to describe ‘rules of thumb’ for realistic sizing.
- Be able to explain the use of the Balanced System Guideline tool.
- Be able to explain the use of the LPAR Validation tool.
- Be able to locate other tools for sizing systems for particular applications.

Figure 9-98. Unit Summary AM711.0

Notes:

Unit 10. Validating the Solution

What This Unit Is About

This unit describes the process of validating a proposed technical solution.

What You Should Be Able to Do

After completing this unit, you should be able to:

- Describe the Solution Assurance Review process
- List the benefits of Solution Assurance
- Determine which type of SAR is appropriate to a particular situation
- Request a Solution Assurance Review
- Conduct a Solution Assurance Review

How You Will Check Your Progress

Accountability:

- Checkpoint

Unit Objectives

On completion of this unit, you should be able to:

- Describe the Solution Assurance Review process.
- List the benefits of Solution Assurance.
- Determine which type of SAR is appropriate to a particular situation.
- Request a Solution Assurance Review.
- Conduct a Solution Assurance Review.

Figure 10-1. Unit Objectives AM711.0

Notes:

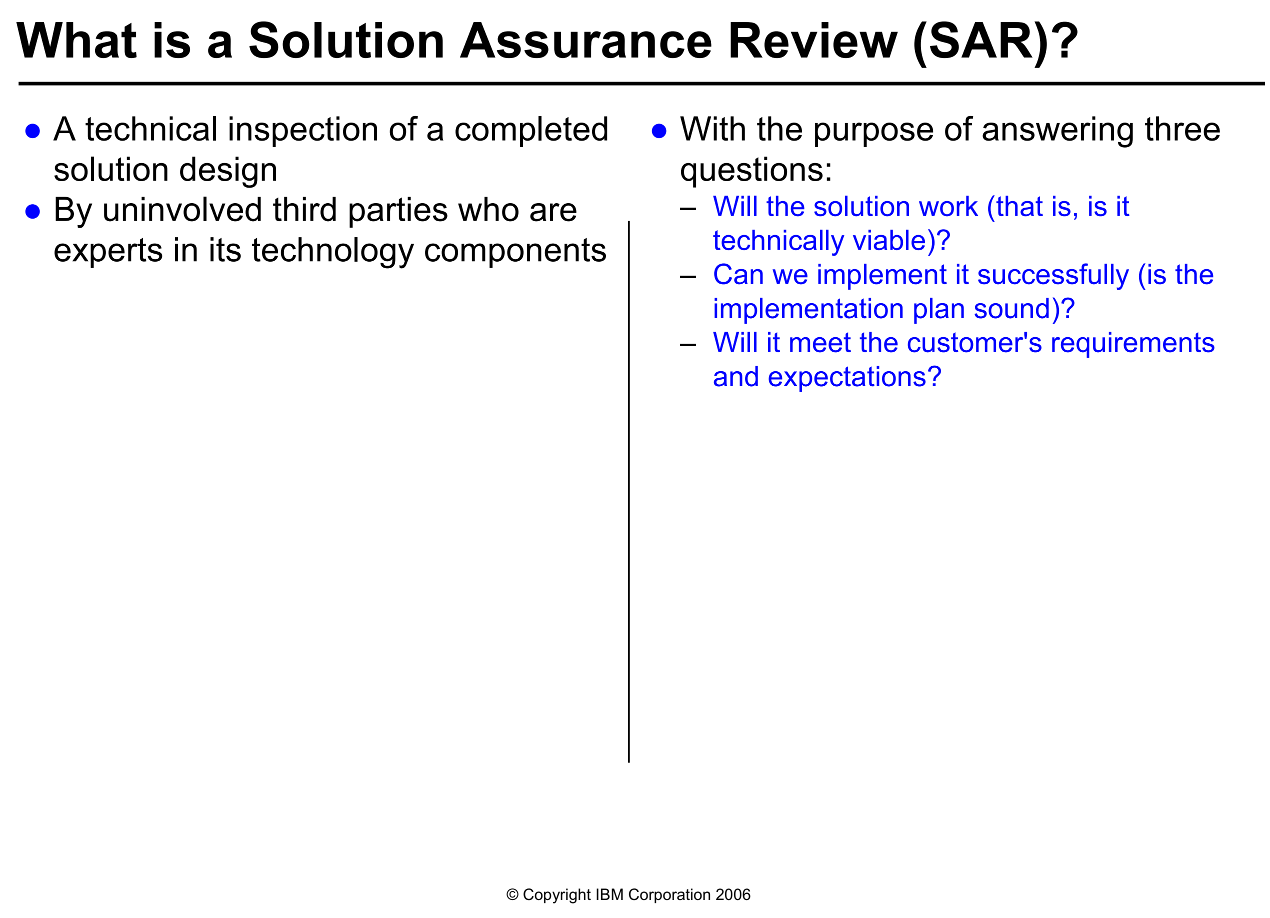

What is a Solution Assurance Review (SAR)?

- A technical inspection of a completed solution design
- By uninvolved third parties who are experts in its technology components
- With the purpose of answering three questions:
  - Will the solution work (that is, is it technically viable)?
  - Can we implement it successfully (is the implementation plan sound)?
  - Will it meet the customer's requirements and expectations?

Figure 10-2. What is a Solution Assurance Review (SAR)? AM711.0

Notes:

- Not all types of SAR require experts external to the proposal team. The details of this are covered later in this section.
- **The Solution Design** process is the series of operations that causes a selected opportunity to be turned into a solution that can be proposed to a customer. As the solution progresses through the design process it is independently reviewed, by way of the Quality Assurance process, at three critical steps that assure that the customer's requirements have been met. These QA steps are:

  Pre-Bid Consulting
  Proposal Baseline Assessment
  Contract Baseline Review

  The most important quality assurance step is the **Solution Assurance Review** as this step establishes that a sound technical design, that will resolve the customer's business problem, can be delivered in a timely manner with the associated risks identified and contained. The **Solution Assurance Review** takes place in the Proposal Baseline

Assessment step of the QA process. Once assured, the solution can then move to the proposal creation and pricing operations.

Depending upon the complexity of a solution, some, or all, of the steps can be combined. Complex solutions, however, should always have separate assurance reviews.

What Are the Benefits of a SAR?

- Increase your selling time
- Minimize installation problems
- Exploit some of IBM’s top technical talent
- Reduce critical situations and their impact on your time
- Set customer expectations realistically and professionally

Figure 10-3. What Are the Benefits of a SAR? AM711.0

Notes:

The Solution Assurance Review process and the documentation generated by it can help both the client and the seller.

Addressing issues that were overlooked during the implementation stage begins can be time consuming, costly, and impair the relationship between the client and the seller.

It is not unusual for there to be problems during the installation and implementation of a solution; SAR can minimize these problems or at least indicate that the problem may occur and allow you to create contingent action plans.

Expert reviews allow your proposed solution to be scrutinized by Subject Matter Experts (SMEs). Their assessment and input can be invaluable in creating a successful solution.

Critical situations demand immediate attention which equals reactive time spent for your team. No matter how the critical situation is resolved it is unlikely that the client will feel more positive toward the relationship with Team IBM than they did before the situation. So you spend a lot of unplanned time just trying to limit damage.

The pre-sale Solution Assurance Review can be a good opportunity to check customer expectations and set them appropriately. Does the client understand the limitations of the solution you are proposing? If they do not, now is the time to address that rather than during the implementation stage.

When Should a SAR Be Done?

- IBM recommends that all solutions be reviewed prior to the proposal being presented to the client.
- The appropriate level of SAR is determined in Stage 4.
- Phase 1: Pre-sale (pre-order)
  - Conducted in Stage 5
- Phase 2: Pre-installation
  - Conducted in Stage 7

Figure 10-4. When Should a SAR Be Done? AM711.0

Notes:

Solution Assurance for pSeries servers should be conducted in two phases:

Phase 1 - Pre-sales (pre-order) review:

IBM recommends that *all* solutions be reviewed prior to the proposal being presented to the client.

Phase 2 - Pre-installation review:

The pre-installation review should be accomplished one to two weeks before the solution/system ships, or earlier if significant porting, moving, or site preparation tasks will be required.

Pre-sale SAR

- Focus on:
  - Capacity planning
  - Sizing
  - Required features

- If Solution involves a migration also consider:
  - Unsupported devices, features, LPPs, and PRPQs
  - Migrated devices and features

Figure 10-5. Pre-sale SAR AM711.0

Notes:

During the pre-order solution review you examine many of the key components of a successful solution, including the results of capacity planning to correctly size server model, required features and, in the case of upgrades, unsupported devices, features, LPPs and PRPQs.

Pre-installation SAR

- Focus on:
  - Power requirements
  - Space requirements
  - Cabling requirements
  - Installation plan and responsibilities
  - Site delivery route considerations
  - Upgrade plan and responsibilities
  - Services and support

Figure 10-6. Pre-installation SAR AM711.0

Notes:

The pre-installation SAR is conducted one or two weeks before the system/solution ships. Allow more time if porting, moving, or site preparation tasks will be required. During this review the following items should be discussed:

Power requirements

Space requirements

Cabling requirements

Installation plan and responsibilities

Site delivery route considerations

Upgrade plan and responsibilities

Services and support

What Type of SAR Should I Perform?

- Type depends on complexity of solution:
  - Self
  - Peer
  - Expert

Figure 10-7. What Type of SAR Should I Perform? AM711.0

Notes:

Self

Performed by a knowledgeable member of the proposal team. Used for proposals of low complexity, or solutions for which there is a great deal of previous implementation experience.

Peer

Performed by a knowledgeable individual on behalf of the proposal team. Used for more complex solutions where there is value in having a reviewer with more subject matter experience, or who has been less involved in the development of the proposal.

Expert

A formal review methodology used for complex solutions or solutions involving new technologies. Outside Subject Matter Experts (SMEs) are used to evaluate the technical details of the proposal. A formal report is generally prepared with recommendations and a risk assessment.

When is an Expert SAR Mandatory?

- p5 590 and 595 systems
- p5 systems implementing any of the following:
  - Logical Partitioning (LPAR)
  - Hardware Management Console (HMC)
  - Linux or i5/OS
  - Systems targeted for Decision Support applications.
- Systems implementing Cluster 1600 with CSM

Figure 10-8. When is an Expert SAR Mandatory? AM711.0

Notes:

590/595

An expert level SAR must be performed at both the pre-sale and pre-installation phases. An IBM Service Representative must participate in the pre-installation SAR called by IBM Business Partners. In addition the "*IBM eServer p5 590/595 Solution Assurance Review (SAR) Confirmation Form*" must be completed. This form can be found at *www.ibm.com/servers/resourcelink.* All 590/595 orders will default to hold status until this documentation is completed.

p5 Systems

The expert SAR is required for the pre-sale phase but the pre-installation SAR can be either self or peer-level.

P5 Systems implementing Cluster 1600 with CSM

For Cluster 1600 solutions an expert review is required for both the pre-sale and pre-installation phases.

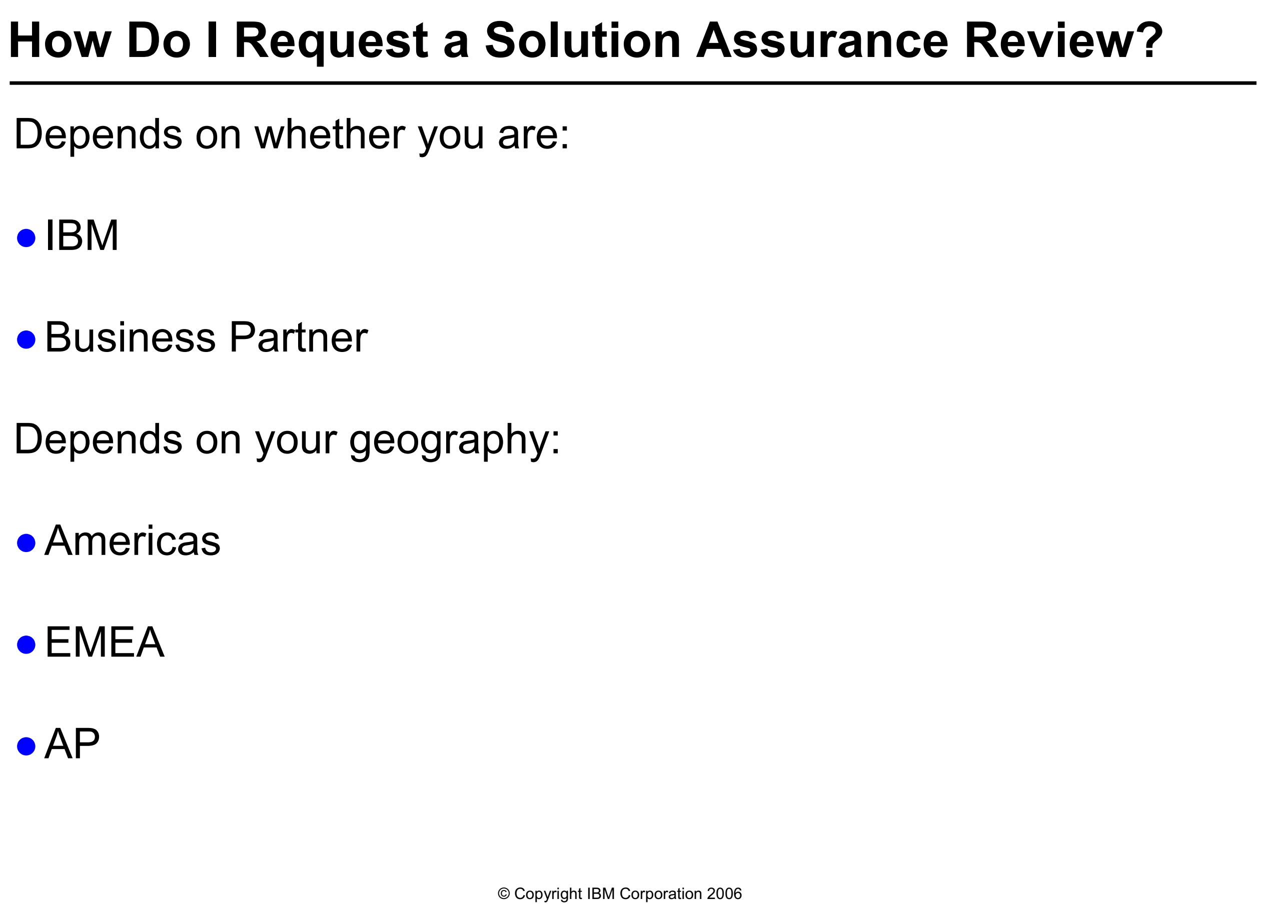

How Do I Request a Solution Assurance Review?

Depends on whether you are:

- IBM
- Business Partner

Depends on your geography:

- Americas
- EMEA
- AP

Figure 10-9. How Do I Request a Solution Assurance Review? AM711.0

Notes:

General Information about Solution Assurance

The IBM Solution Assurance Web site

IBM http://w3-03.ibm.com/support/assure/assur30i.nsf/Web/SA

Business Partner
http://www-.ibm.com/partnerworld/pwhome.nsf/weblook/trs_wgd_sa.html

To request an SAR or for assistance:

Americas IBM sales representatives may request a solution assurance for pSeries servers they sell via TechXpress. To request assistance, go to: w3.ibm.com/support -> Request Technical Sales resources - Americas, or call the Sales Productivity Center at 877-707-2727 for the US and Canada; or 770-858-5451 for Latin America.

EMEA: To request a Solution Assurance review or for assistance:

IBM - refer to this link: w3-1.ibm.com/support/assure/assur30i.nsf/WebIndex/SA424 *or*

IBM Business Partners - contact technical sales support in your country

AP: To request a Solutions Assurance review or for assistance, refer to this link: w3-1.ibm.com/support/assure/assur30i.nsf/WebIndex/SA441

The Client Environment and Expectations

- Requirements
- Conditions of satisfaction
- Assumptions and risks
- Project dependencies
- Sponsorship
- Technical environment
- Client understanding
- Client escalation history
- Current issues
- Services planned
- Other pSeries systems
- Staffing and skill levels

Figure 10-10. The Client Environment and Expectations AM711.0

Notes:

More detailed information about the client environment can be found in the Solution Assurance Discussion Guide which the instructor will provide you with.

Pre-Sales Checklists

- Capacity/Performance/Sizing
- LPAR
- Availability
- Disk Storage
- Hardware Management Console
- System Management Environment
- General Hardware Considerations

Figure 10-11. Pre-Sales Checklists AM711.0

Notes:

More detailed information about the Pre-Sales Checklists can be found in the Solution Assurance Discussion Guide which the instructor will provide you with.

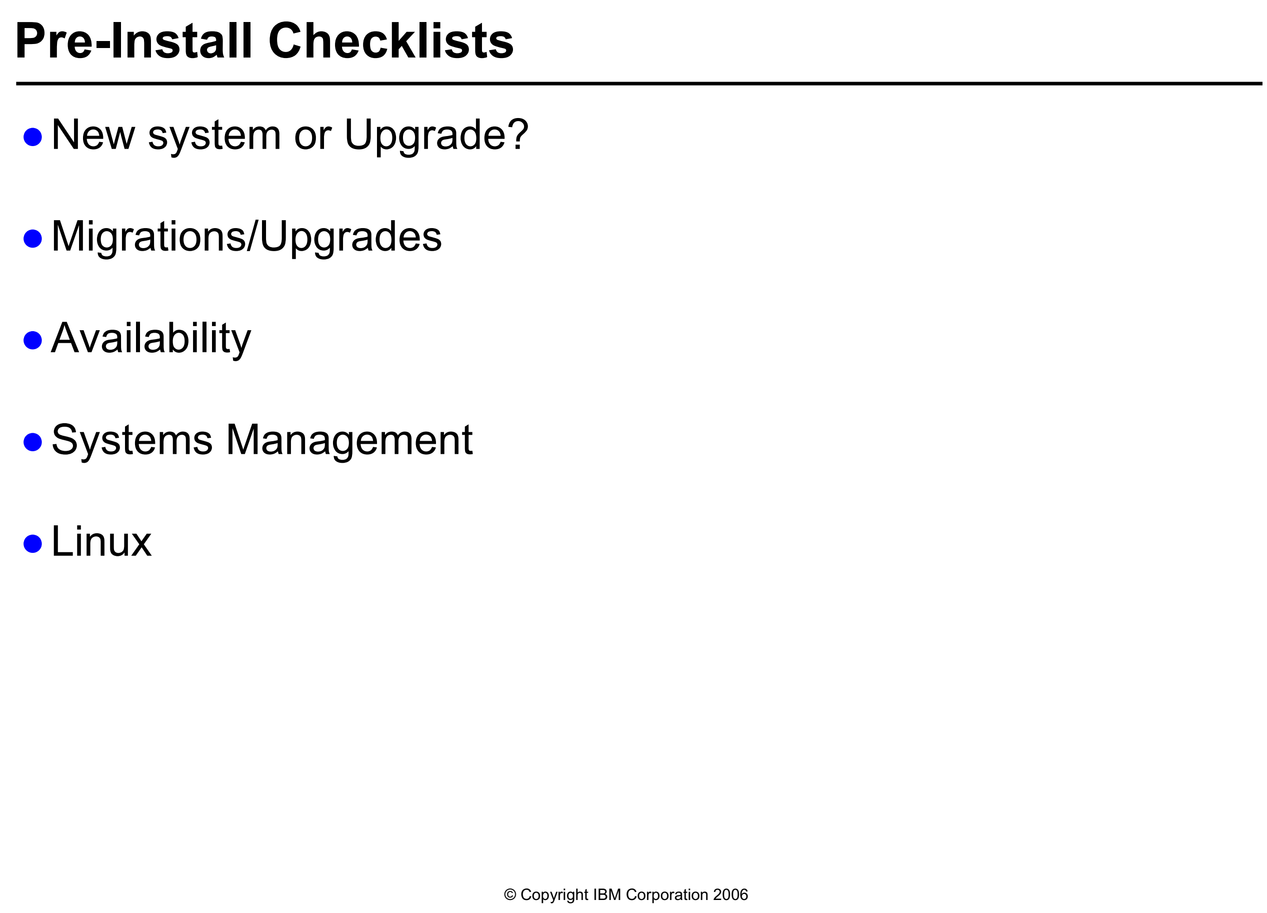

Pre-Install Checklists

- New system or Upgrade?
- Migrations/Upgrades
- Availability
- Systems Management
- Linux

Figure 10-12. Pre-Install Checklists AM711.0

Notes:

More detailed information about the Pre-Install Checklists can be found in the Solution Assurance Discussion Guide which the instructor will provide you with.

Product Specific Checklists

- p5 590 and p5 595
- HMC
- Operating Systems
- Virtualization/LPAR
- Clustering
- Capacity on Demand (CoD)

Figure 10-13. Product Specific Checklists AM711.0

Notes:

More detailed information about the Product Specific Checklists can be found in the Solution Assurance Discussion Guide which the instructor will provide you with.

Checkpoint Questions (1 of 2)

1. A Solution Assurance Review is:
   a. An opportunity to add new functionality to the solution prior to presenting the proposal to the client.
   b. A technical inspection of a completed solution design.
   c. The document signed off by the client on the day the solution is delivered.
   d. An opportunity to provide training for the client's system administrators.

2. The benefits of Solution Assurance Review are:
   a. To minimize installation problems
   b. To reduce critical situations and their impact on your time.
   c. To set customer expectations realistically and professionally
   d. All of the above

Figure 10-14. Checkpoint Questions (1 of 2) AM711.0

Notes:

Checkpoint Questions (2 of 2)

3. The three types of SAR are:
   a. Internal, external and mixed
   b. Self, peer, and expert
   c. Quick, detailed, and online
   d. Mandatory, optional, and online

4. Which of the following topics are covered in the pre-installation SAR?
   a. Power, space and cabling requirements
   b. Installation plan and responsibilities
   c. Services and support
   d. All of the above

5. True or False. Any proposed solution which includes the p5 590 or 595 system must have an expert SAR.

Figure 10-15. Checkpoint Questions (2 of 2) AM711.0

Notes:

Any Questions?

Figure 10-16. Any Questions? AM711.0

Notes:

Unit Summary

You should now be able to:

- Describe the Solution Assurance Review process.
- List the benefits of Solution Assurance.
- Determine which type of SAR is appropriate to a particular situation.
- Request a Solution Assurance Review.
- Conduct a Solution Assurance Review.

Figure 10-17. Unit Summary AM711.0

Notes:

References

- You may find more information here:
  The location of the most recent SA Discussion Guide

  http://w3.ibm.com/support/assure/ Search for document number SA675)
  The eSAR Home Page
  http://www-1.ibm.com/esar/esarhome.nsf/hps?OpenForm&d=eSARHome
  The Solution Assurance Website

  http://w3-03.ibm.com/support/assure/assur30i.nsf/Web/SA

Figure 10-18. References AM711.0

Notes:

Depending on your geography and whether or not you are an IBMer or Business Partner these URLs should provide you with access to more information about Solution Assurance.

Let's take a break

Figure 10-19. Let's take a break AM711.0

Notes:

Appendix A. Checkpoint Solutions

Unit 2: The Sales Cycle

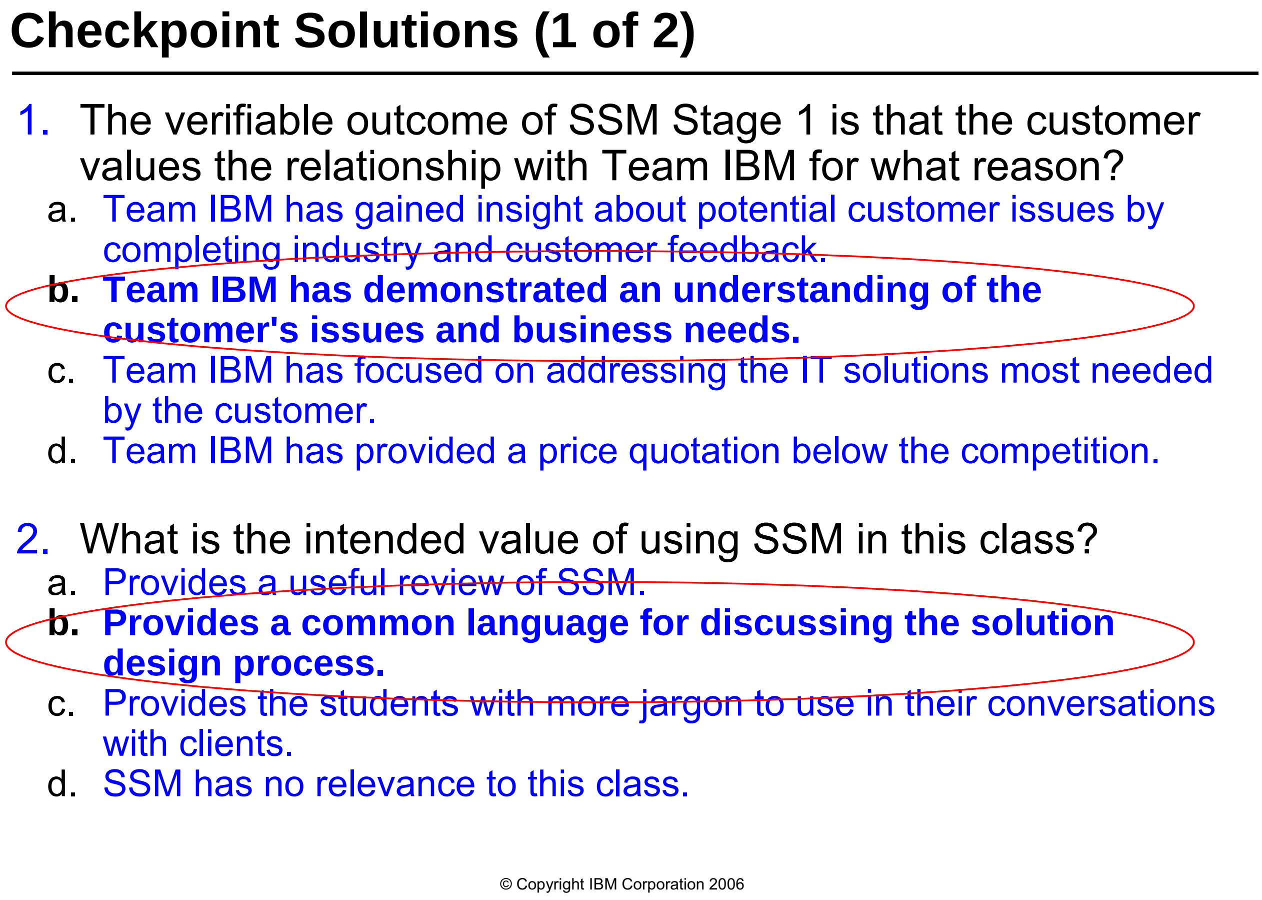

Checkpoint Solutions (1 of 2)

1. The verifiable outcome of SSM Stage 1 is that the customer values the relationship with Team IBM for what reason?
   a. Team IBM has gained insight about potential customer issues by completing industry and customer feedback.
   [✓ b. **Team IBM has demonstrated an understanding of the customer's issues and business needs.**]
   c. Team IBM has focused on addressing the IT solutions most needed by the customer.
   d. Team IBM has provided a price quotation below the competition.

2. What is the intended value of using SSM in this class?
   a. Provides a useful review of SSM.
   [✓ b. **Provides a common language for discussing the solution design process.**]
   c. Provides the students with more jargon to use in their conversations with clients.
   d. SSM has no relevance to this class.

Checkpoint Solutions (2 of 2)

1. At what Stage is a pre-sale Solution Assurance Review performed?
   a. Stage 1
   b. Stage 3
   c. Stage 6
   [✓ **d. Stage 5**]

2. In SSM terms a person in the customer's organization who typically believes that a seller's success is critical to the organization or to that person is called:
   a. A decision maker
   [✓ **b. A sponsor**]
   c. An opportunity magnet
   d. A decision leader

Unit 3: Technology

Checkpoint Solutions (1 of 7)

1. [✓ **True**] or False. RISC architecture was invented by IBM.

2. The acronym RISC stands for:
   - [✓ a. **Reduced Instruction Set Computer**]
   - b. Really Interesting System Concept
   - c. Radically Intense Silly Computer
   - d. None of the above

3. [✓ **True**] or False. The acronym POWER stands for Performance Optimization with Enhanced RISC.

4. In measuring execution time of a program which of the following factors are important?
   - a. Cycle time
   - b. Compiler optimization
   - c. Number of instructions/cycle
   - [✓ d. **All of the above are important factors**]

Checkpoint Solutions (2 of 7)

5. Which System p feature helps eliminate intermittent errors?
   a. Service processor
   b. Service agent
   [✓ c. **First failure data capture**]
   d. Dynamic processor deallocation

6. Which of the following is not related to System p RAS?
   a. Light path diagnostics
   b. First failure data capture
   [✓ c. **Simultaneous multi-threading**]
   d. Dynamic processor de-allocation

Checkpoint Solutions (3 of 7)

7. Which of the following does not provide additional reliability for System p memory?
   a. Bit-scattering
   b. Bit-steering
   c. Bit-scrubbing
   [✓ **d. Bit-of-this-bit-of-that**]

8. What is true about System p processors?
   a. They can be hot-swapped for convenience
   b. They never fail due to first failure data capture
   [✓ **c. They may be implemented as hot spares**]
   d. They can be fixed via firmware

Checkpoint Solutions (4 of 7)

9. Partitions which have limited access to processor resources in the shared pool are____________.
    a. Uncapped
    [✓ **b. Capped**]
    c. Virtual
    d. None of the above

10. [✓ **True**] or False. When more than one uncapped partition is competing for processor resources in the shared pool the preference is given to the partition with the highest priority.

11. [✓ **True**] or False. If a partition is unable to acquire the minimum processor resources required it will not activate.

Checkpoint Solutions (5 of 7)

12. [✓ **True**] or False. Virtual I/O allows the provision of disk and Ethernet devices without actually having physical adapters in the partition.

13. True or [✓ **False**]. Virtual SCSI disks can be created on POWER4 systems.

14. [✓ **True**] or False. A virtual SCSI client requires a virtual SCSI server.

15. [✓ **True**] or False. Virtual Ethernet connections between partitions do not require a physical network interface card in either partition.

16. True or [✓ **False**]. A virtual SCSI disk cannot be provided by a physical fibre channel disk.

17. [✓ **True**] or False. System boot disks are good candidates for virtual disks.

Checkpoint Solutions (6 of 7)

18. [✓ **True**] or False. Reserve Capacity on Demand allows activation of processor resources in a shared pool.

19. [✓ **True**] or False. On/Off Capacity on Demand is a post pay billable approach to Capacity on Demand which is available for both POWER4 and POWER5 systems.

20. To accommodate a client’s planned growth in processing demand which of the following options would be suggested.
    a. [✓ **Permanent Capacity on Demand**]
    b. Reserve Capacity on Demand
    c. On/Off Capacity on Demand
    d. Trial Capacity on Demand

21. True or [✓ **False.**] Trial Capacity on Demand can be activated multiple times provided the total time does not exceed the maximum entitlement.

Checkpoint Solutions (7 of 7)

22. True or [✓ **False**]. Any technology improvement will boost performance of any client solution.

23. [✓ **True**] or False. The application of technology in a creative way to solve clients’ business problems is one definition of innovation.

24. A client’s satisfaction with your solution can be enhanced by which of the following?
    a. Setting expectations appropriately.
    b. Applying technology appropriately.
    c. Communicating the benefits of the technology to the client.
    [✓ **d. All of the above.**]

25. Which of the following are available with POWER5 architecture?
    a. Simultaneous Multi-Threading
    b. Micro-Partitioning
    c. Dynamic power management
    [✓ **d. All of the above**]

26. True or [✓ **False**]. Simultaneous multi-threading is the same as hyperthreading, IBM just gave it a different name.

Unit 4: Servers

Checkpoint Solutions (1 of 4)

1. Which of the following are attributes of the IBM eServer brand?
   [✓ **a. Broad range of products**]
   b. Advanced RAS features
   c. Linux across the whole product range
   d. All of the above

2. IBM eServer pSeries servers are_____________.
   a. Able to run AIX and Windows XP
   [✓ **b. Able to run AIX and Linux**]
   c. Able to run only AIX
   d. Able to run only Linux

3. How many 7311-D20 I/O drawers can be attached to the IBM System p5 560Q?
   a. 8
   b. 20
   c. 4
   [✓ **d. No I/O drawers can be attached to this system**]

Checkpoint Solutions (2 of 4)

4. Which of the following is true of the System p5 575 server.
   - a. It can support 8 or 16 CPUs
   - b. It can support clock speeds of 1.5 GHz only
   - c. It can support up to 256 GB of RAM
   - [✓ **d. Both a and c**]
5. A client wants a deskside workstation that can support Altivec extensions and 3D graphics. Which of the following correctly describes the options available?
   - a. Only the p5-550 can provide this.
   - b. There are no deskside workstations that can support these requirements.
   - c. Either the p5-550 or the p5-520 can support these requirements.
   - [✓ **d. Only the IntelliStation POWER 185 Express workstation can provide this**.]
6. Which of the following are true statements about the p5-595 server?
   - a. From 16 to 64 processors are supported.
   - b. A maximum of 254 micropartitions are supported.
   - c. 8 GB of RAM is standard.
   - [✓ **d. All of the above.**]

Checkpoint Solutions (3 of 4)

7. Which of the following statements is true about the IBM System p5 570 server?
   - a. It can support processor speeds of 1.9 GHz and 2.2 GHz
   - b. It can support up to 16 processor cores per system
   - c. It can support up to 20 I/O drawers
   - [✓ **d. All of the above.**]
8. Which server offers quad-core technology in a 2U package?
   - [✓ **a. p5-510Q**]
   - b. p5-550Q
   - c. p5-570Q
   - d. No current system provides this capability.

Checkpoint Solutions (4 of 4)

9. Which of the following statements is true about the JS20 Blade.
    - [✓ **a. It can run both AIX and Linux**.]
    - b. It can support micropartitions.
    - c. It can support SCSI disks internally.
    - d. All of the above.

10. The D20 I/O drawer provides support for:
    - [✓ **a. 12 hotplug SCSI disks and 7 PCI-X slots**]
    - b. 6 hotplug SCSI disks and 6 PCI-X slots
    - c. 14 PCI-X slots
    - d. None of the above

Unit 5: Infrastructure

Checkpoint Solutions (1 of 6)

1. Which of the following is not a benefit of networking?
   a. Shared access to resources
   b. Collaborative computing
   c. Simplification of connection between technologies
   [✓ **d. Increased system security**]

2. Which of the following is not true about System p?
   2. Support for many network adapter types
   [✓ **3. Support only for IBM proprietary adapters**]
   4. Operating systems support for advanced networking features
   5. Ease of integration into existing network infrastructure

3. True or [✓ **False**]. A single copper cable segment can carry Ethernet traffic over 20 kilometers.

4. [✓ **True**] or False. TCP/IP is a suite of protocols.

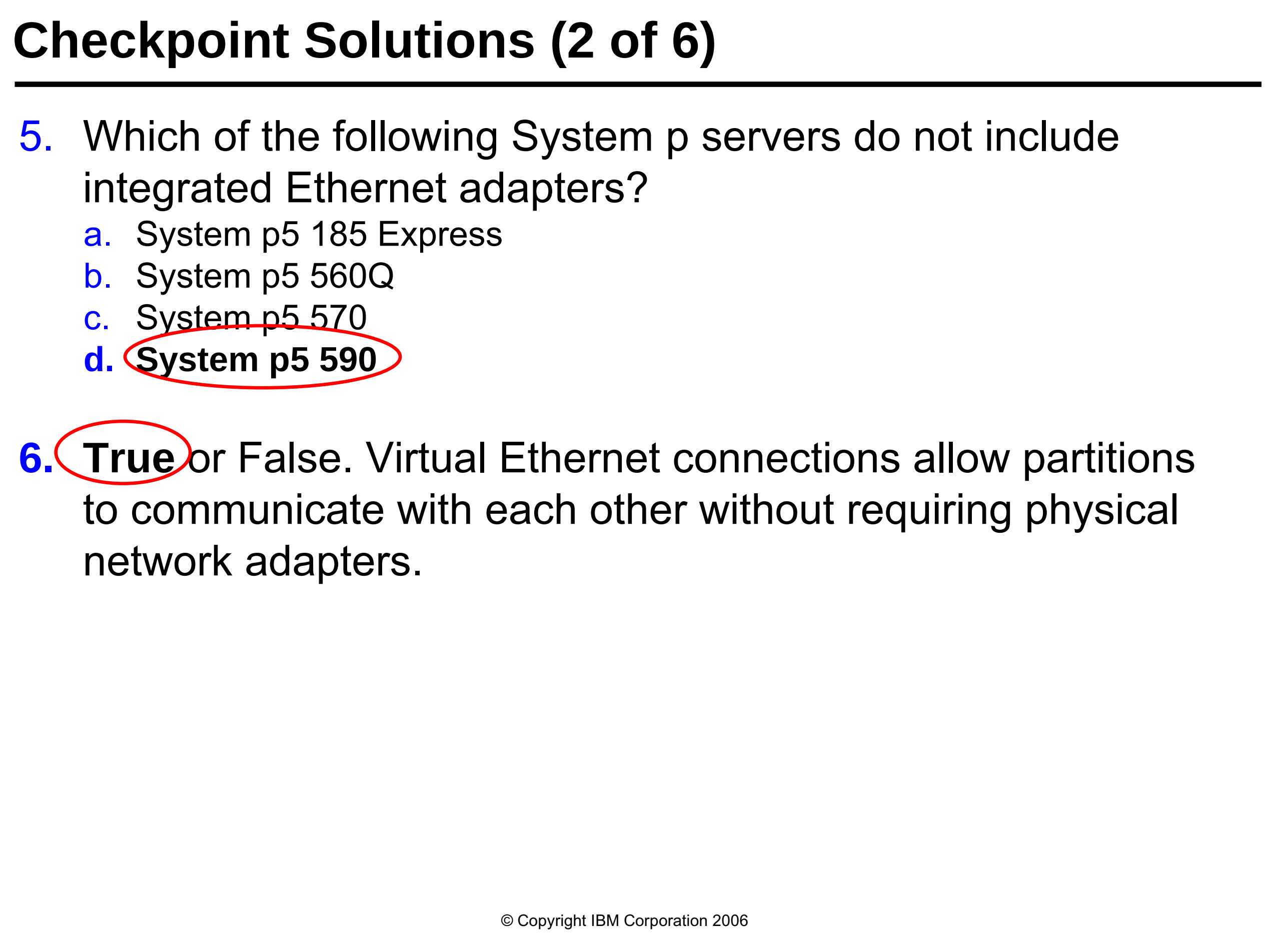

Checkpoint Solutions (2 of 6)

5. Which of the following System p servers do not include integrated Ethernet adapters?
   a. System p5 185 Express
   b. System p5 560Q
   c. System p5 570
   d. [✓ System p5 590]

6. [✓ **True**] or False. Virtual Ethernet connections allow partitions to communicate with each other without requiring physical network adapters.

Checkpoint Solutions (3 of 6)

7. Which of the following is not a benefit of clustering?
   a. Enhanced scalability
   b. Application availability
   c. Simplified system management
   [✓ **d. Fault tolerance**]

8. True or [✓ **False**]. Cluster 1600 systems are only used for high-performance computing in technical and scientific environments.

9. [✓ **True**] or False. A cluster node can be a server, a blade, or a logical partition.

10. Which of the following methods can be used to connect cluster nodes together?
    a. High performance switch
    b. Myrinet-2000 switch
    c. InfiniBand switch
    [✓ **d. All of the above**]

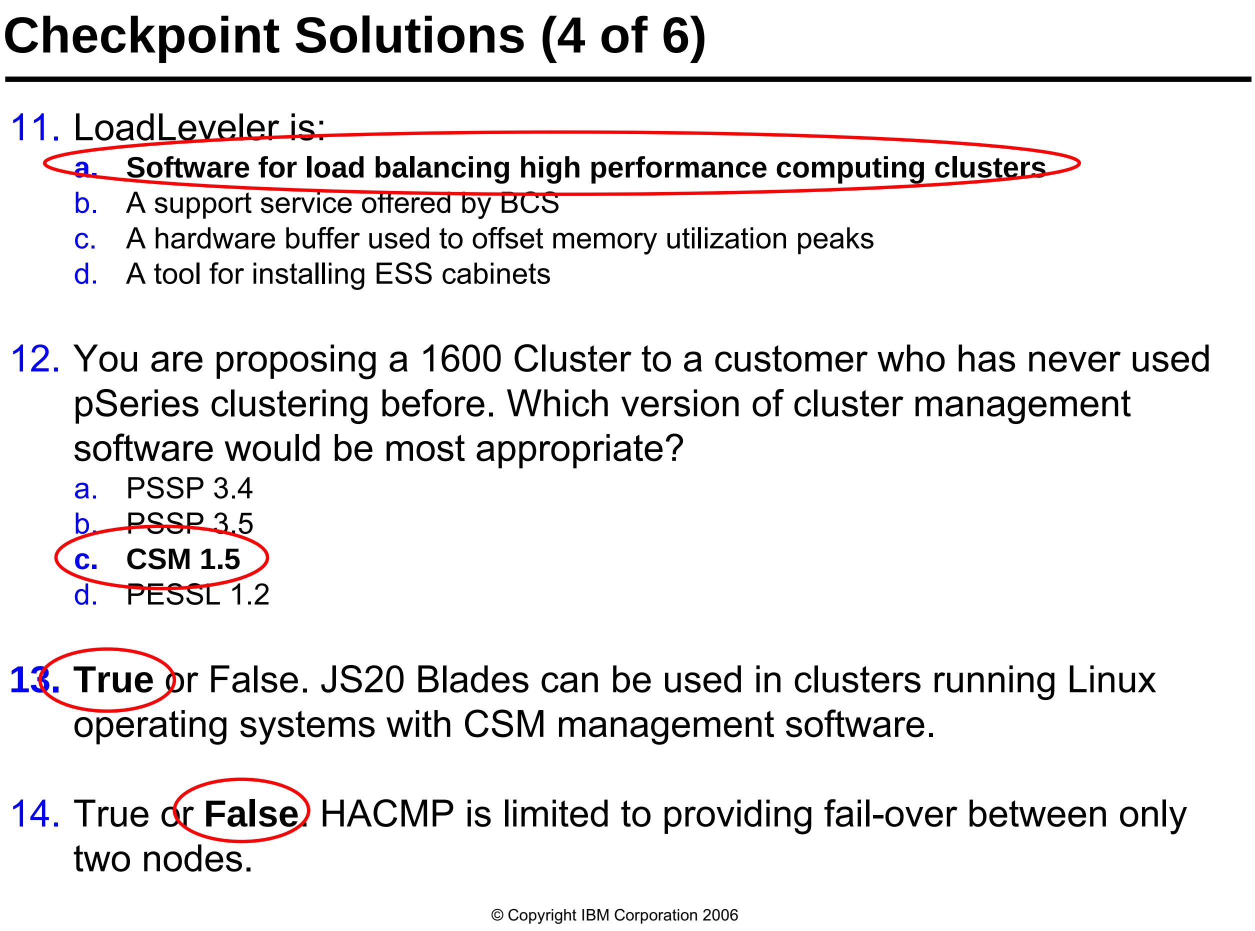

Checkpoint Solutions (4 of 6)

11. LoadLeveler is:
    - [✓ a. **Software for load balancing high performance computing clusters**]
    - b. A support service offered by BCS
    - c. A hardware buffer used to offset memory utilization peaks
    - d. A tool for installing ESS cabinets

12. You are proposing a 1600 Cluster to a customer who has never used pSeries clustering before. Which version of cluster management software would be most appropriate?
    - a. PSSP 3.4
    - b. PSSP 3.5
    - [✓ c. **CSM 1.5**]
    - d. PESSL 1.2

13. [✓ **True**] or False. JS20 Blades can be used in clusters running Linux operating systems with CSM management software.

14. True or [✓ **False**]. HACMP is limited to providing fail-over between only two nodes.

Checkpoint Solutions (5 of 6)

15. What is Grid computing?
    - a. Distributed computing based on open standards for enabling heterogeneous resources.
    - b. A computing model that optimizes the productivity of your IT resources by maximizing utilization levels without a considerable increase in operating cost.
    - c. A critical component of IBM’s On Demand vision.
    - [✓ **d. All of the above.**]

16. One obstacle to enterprise-wide Grid computing is:
    - [✓ **a. Political issues between different lines of business.**]
    - b. The technology to support this is not available.
    - c. There are not sufficient accounting tools available.
    - d. A lack of scalability of the Grid computing model.

Checkpoint Solutions (6 of 6)

17. Which of the following are benefits of Grid computing?
    a. Application Acceleration
    b. Aggregating Information
    c. Infrastructure Flexibility
    [✓ **d. All of the above**]

18. Which of the following is a Grid computing focus area?
    a. Business Analytics
    b. Engineering and Design
    c. Enterprise Optimization
    [✓ **d. All of the above**]

Unit 6: Storage

Checkpoint Solutions (1 of 2)

1. IBM System Storage offerings:
   a. Benefit from the synergy between Storage and Server development.
   b. Are a part of the On Demand operating environment.
   c. Contribute to the performance achievements of System p servers.
   [✓ d. **All of the above**]

2. IBM System Storage offerings:
   a. Include management software as well as hardware.
   b. Cover the range from entry-level to enterprise requirements.
   c. Can interoperate with other vendors’ servers.
   [✓ d. **All of the above**]

Checkpoint Solutions (2 of 2)

3. The SAN Volume Controller:
   - a. Combines storage capacity into a single storage pool.
   - b. Allows management of the storage pool from a central point.
   - c. Reduces down time.
   - [✓ **d. All of the above**]

4. True or [✓ **False**]. SAN and NAS are simply different names for the same storage concept.

5. LTO Ultrium 3 tapes have a compressed storage capacity of ______.
   - a. 100 Gigabytes
   - b. 200 Gigabytes
   - [✓ **c. 800 Gigabytes**]
   - d. 1000 Gigabytes

Unit 7: Software

Checkpoint Solutions (1 of 6)

1. Which of the following are true statements about AIX 5L as a market leader among UNIX operating systems?
   - a) Mainframe-inspired capabilities
   - b) Broad application support
   - c) Available on Intel and Power architecture
   - [✓ **d) Statements a and b are true**]

2. The client is interested in being able to move IT resources dynamically. Which of the following advantages of AIX on Power5 architecture would you emphasize?
   - a. Reliability, availability, and serviceability.
   - [✓ **b. Flexible resource model**]
   - c. Network performance
   - d. Systems management

Checkpoint Solutions (2 of 6)

3. Currently the UNIX marketplace is characterized by:
   a. Rapid growth and the emergence of numbers of small vendors.
   b. A rapid decline as Windows operating systems dominate all segments.
   [✓ **c. Flat growth or small declines with fierce competition for market share**]
   d. Rapid growth of Solaris

4. Which of the following is feature introduced at AIX5L V5.3 TL5?
   a. JFS2 file system support
   [✓ **b. Shared boot image management**]
   c. NFSV3 support
   d. KDE window manager support

Checkpoint Solutions (3 of 6)

5. When proposing a POWER Linux solution, the features to highlight are:
   - [✓ **a. RAS, Virtualization, range of servers, roadmap**]
   - b. SUSE, Virtualization, Linux affinity
   - c. Ordering of Linux with POWER systems, support structure, FTSS support
   - d. Low environmental costs, large solution portfolio, AIX and Linux interchangeability

6. The following is TRUE for Linux on POWER:
   - a. All POWER-based servers are supported consistently and completely by SuSE 8 Linux
   - b. All POWER servers are supported consistently by all Linux distributions
   - [✓ **c. Different POWER servers have different levels of Linux distribution support**]
   - d. All POWER servers are consistently supported by the IBM Linux distribution

Checkpoint Solutions (4 of 6)

7. The POWER-based Linux solutions are particularly attractive to:
   a. Customers with a single dedicated server firewall requirement.
   b. Customers who need extensive vertical application support in the SMB environment with Windows applications.
   [✓ **c. Customers who plan to develop and implement solutions requiring high performance, high reliability, and at a surprisingly low price, starting with the JS20 and OpenPower 710.**]
   d. Customers who implemented Linux on IBM System z.

8. [✓ **True**] or False. Most Linux on IBM platforms run on Intel.

9. True or [✓ **False**]. There are no functional benefits from running Linux on Power versus Linux on Intel.

Checkpoint Solutions (5 of 6)

10. Solution offerings add value because they are tested and sized by IBM and include at least three (SML) recommended configurations.
    - [✓ **a. True**]
    - b. False

11. Reduced complexity, lower price/performance and the ability to run more workloads on fewer servers are benefits of server consolidation.
    - [✓ **a. True**]
    - b. False

12. The IBM System p solution offering consists entirely of IBM products and services.
    - a. True
    - [✓ **b. False**]

Checkpoint Solutions (6 of 6)

13. Propose BladeCenter JS21 as a Linux web server solution because:
    a. You can capture Sun and HP clients ready to implement Web serving on Linux.
    b. Infrastructure is the first area clients use Linux.
    c. BladeCenter JS21 strengths play to the client infrastructure needs.
    [✓ **d. All of the above**]

14. Which of the following are signs of an IBM Global Financing opportunity?
    a. Projects with small upfront costs
    b. Minimal license fee and/or application development investment
    [✓ **c. Clients who want to match cost outlays to business case benefits**]
    d. All of the above

Unit 8: Competition

Checkpoint Solutions (1 of 2)

1. True or [✓ **False.**] Commercial benchmarks results can be easily compared to the client’s workload.

2. IBM System p has a competitive advantage in which of the following?
   a. Performance
   b. RAS
   c. Virtualization
   [✓ **d. All of the above**]

3. Which of the following is an advantage of logical partitioning over physical partitioning?
   a. Granularity
   b. Resource sharing
   c. Flexibility
   [✓ **d. All of the above**]

Checkpoint Solutions (2 of 2)

4. The System p5-505 is competitive against the Sun X4100 in which of the following:
   a. Runs Windows as well as Linux and AIX
   b. Price/performance
   c. High density
   [✓ d. **Both b and c**]

5. The System p5-520Q competes with the Sun V40z in which areas?
   a. I/O slots
   b. Drive bay capacity
   c. Price/performance
   [✓ d. **All of the above**]

Checkpoint Solutions (1 of 2)

1. Usually the main obstacle to accurate sizing is:
   a. Inaccurate tools
   b. Too much information
   [✓ **c. Not enough information**]
   d. Lack of interest in accurate sizing

2. The rules of thumb for sizing are based on:
   a. Measurements taken on the client’s system
   [✓ **b. Experience**]
   c. Pure guesswork
   d. Astrology

Checkpoint Solutions (2 of 2)

3. Which of the following questions is useful in segmenting a sizing problem?
   a. What is the type of application?
   b. What is the user doing? (workload type)
   c. How much data is being processed?
   [✓ **d. All of the above**]

4. If the run queue length is greater than 10 times the number of processors in the system then, from a utilization point of view this is:
   a. Good
   b. Bad
   [✓ **c. Ugly**]
   d. Normal

Checkpoint Solutions (1 of 2)

1. A Solution Assurance Review is:
   a. An opportunity to add new functionality to the solution prior to presenting the proposal to the client.
   [✓ **b. A technical inspection of a completed solution design.**]
   c. The document signed off by the client on the day the solution is delivered.
   d. An opportunity to provide training for the client's system administrators.

2. The benefits of Solution Assurance Review are:
   a. To minimize installation problems
   b. To reduce critical situations and their impact on your time.
   c. To set customer expectations realistically and professionally
   [✓ **d. All of the above**]

Checkpoint Solutions (2 of 2)

3. The three types of SAR are:
   - a. Internal, external and mixed
   - [✓ **b. Self, peer, and expert**]
   - c. Quick, detailed, and online
   - d. Mandatory, optional, and online

4. Which of the following topics are covered in the pre-installation SAR?
   - a. Power, space and cabling requirements
   - b. Installation plan and responsibilities
   - c. Services and support
   - [✓ **d. All of the above**]

5. [✓ **True**] or False. Any proposed solution which includes the p5 590 or 595 system must have an expert SAR.